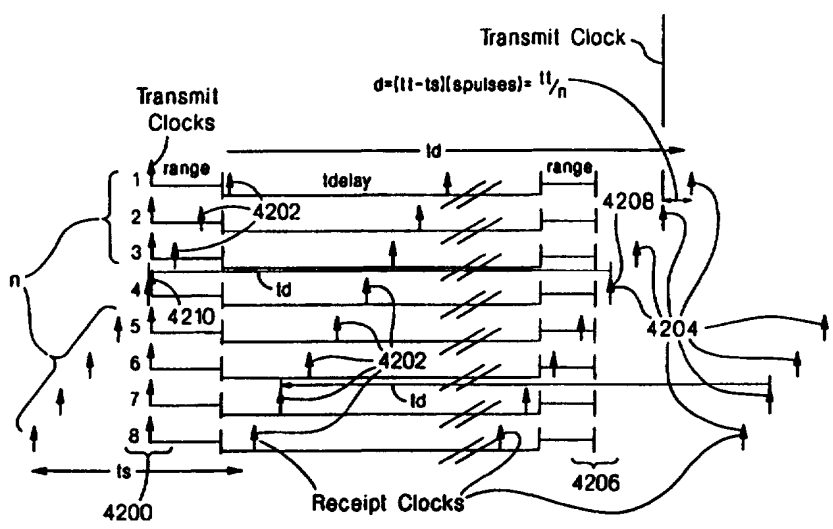




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification 6 : G01C</p>	<p>A2</p>	<p>(11) International Publication Number: WO 97/49972 (43) International Publication Date: 31 December 1997 (31.12.97)</p>
<p>(21) International Application Number: PCT/US97/11249 (22) International Filing Date: 26 June 1997 (26.06.97) (30) Priority Data: 08/670,612 26 June 1996 (26.06.96) US (71) Applicant: PAR GOVERNMENT SYSTEMS CORPORATION [US/US]; 8383 Seneca Turnpike, New Hartford, NY 13413-4991 (US). (72) Inventor: WOOLLEY, Louis, A.; 3686 Fountain Street, Clinton, NY 13323 (US). (74) Agent: WALPERT, Gary, A.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).</p>	<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>Without international search report and to be republished upon receipt of that report.</i></p>	

(54) Title: MEASURING DISTANCE



(57) Abstract

An object in a storage area or moving vehicle is monitored by attaching an electronic tag to the object. An electronic device detects the presence of the object by communicating with the tag while the object is in storage or is being moved by the vehicle. The tags may also determine the location of an attached object and may reroute the object if it deviates from a given shipping schedule. A group of objects is monitored by two electronic tags, each attached to an object in the group. Each tag has circuitry for communicating information relating to an object in the group to a second tag. Each tag also includes a memory connected to the circuitry that is capable of storing the information, and a controller connected to the memory and the circuitry. A distance is measured by transmitting multiple symbols from one object to another object, having the symbols returned such that the symbols' measured round-trip times are not all identical, and calculating the distance using the measured round-trip times.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

- 1 -

MEASURING DISTANCEBackground of the Invention

This invention relates to managing assets equipped
5 with active electronic tags.

Asset management systems attempt to monitor the
existence, location and status of items being transported
or stored. The systems include two main types: barrier
systems and continuous systems.

10 Barrier systems determine whether an asset is
located within a cell delimited by a set of stationary
interrogators. The interrogators detect the asset as it
crosses a boundary of the cell, but are incapable of
determining a position of the asset within the cell.

15 Typical barrier systems fix a bar code to each
asset, and use bar code readers as interrogators.
Examples include systems used by car factories and mail
carriers, such as Federal Express and United Parcel
Service.

20 Another type of barrier system attaches a
transponder to each asset and uses radio transceivers as
interrogators. The transponders are passive, that is,
they respond when activated by an interrogator but are
incapable of initiating communications or other
25 functions. An example includes a railroad system
developed by AMTECH, in which each railroad car on a
train is equipped with a transponder. Interrogators
include transceivers that are placed along a railway.
Continuous systems determine the position of an asset
30 anywhere within a monitoring area. An example,
implemented by LOJACK, equips cars with an active
transponder. The transponder is active because it
initiates its own functions, for example, it monitors a
particular hailing frequency. If the car is stolen,
35 LOJACK uses special transmitters to establish
communication with the transponder at the hailing

- 2 -

frequency. Receivers provided to police cars analyze the communications to locate the car.

Summary of the Invention

In general, in one aspect, the invention features
5 a method for monitoring an object in a moving vehicle by attaching an electronic tag to the object. An electronic device detects the presence of the object by communicating with the tag while the object is being moved by the vehicle.

10 Preferred embodiments of this aspect of the invention include the following features.

The absence of the object is reported to a user or a computer accessible to the user when the device is unable to communicate with the tag. The absence may be
15 reported periodically to either an owner of the object or a possessor of the object.

A location of the object is determined by establishing communications between the tag and a sensor storing location information. The location of the object
20 is compared to a planned route, and any discrepancy between the location and the planned route is reported to the user. The discrepancy may be reported through a computer. The user determines a new route for the object and notifies the elective device of the new route. The
25 user may determine the new route using a computer.

The object may be disposed in a container to which is attached an electronic tag. Communication is established between the tag on the object and the tag on the container. The tag on the object communicates with
30 the electronic device through the tag on the container. The container may also be placed in a second container to which is attached an electronic tag. The tag on the first container communicates with the tag on the second container.

- 3 -

Each tag stores a unique identification for the object, a description of the object, the object's shipper and destination and a signature of each possessor of the object. The signatures stored by the tag are forwarded
5 to the electronic device when the tag reaches the limit of its memory capacity or when the object to which the tag is attached reaches the end of its monitoring period. The signatures are then deleted from the tag memory.

In general, in another aspect, the invention
10 features a system for monitoring and affecting actions on behalf of a group of objects. A first electronic tag attached to an object includes circuitry for communicating with a second electronics tag, the circuitry communicating at least information pertaining
15 to an object in the group. The first tag also includes a memory connected to the circuitry and capable of storing the information. The second electronic tag attached to another object includes circuitry for communicating with the first tag and a memory connected in an analogous
20 manner.

Preferred embodiments of this aspect of the invention include the following features.

The first and second tags each include a controller connected to the memory and the communication
25 circuitry. The information communicated between the tags includes a distance between the tags. The first tag has a time measuring device connected to the controller, the device measuring a time duration required for a signal to travel between the first tag and the second
30 tag. The device in the first tag includes a clocked counter which has a first count corresponding to the time of transmission of the signal and which subsequently has a second count corresponding to the time of recognition that the signal has been received. The device uses the

- 4 -

product of the clock period and the difference between the counts to determine the time duration.

The second tag may also include a delay circuit connected between the circuitry and the controller, the
5 delay circuit causing the second tag to delay returning the signal to the first tag by a given time period. The memory in the first tag includes a program causing the controller to execute a function when the controller reads the memory, the function causing the controller to
10 determine the distance based on the time duration. The second tag includes a second counter connected to the controller in the second tag. The counter in the first tag has a mechanism for adjusting a frequency of the first-mentioned counter, the mechanism being coupled to
15 the controller in the first tag. The program in the first tag causes the controller to employ the mechanism to offset the frequency of the first-mentioned counter from a frequency of the counter in the second tag based on the frequency of the counter in the second tag.

20 The program in the memory of the first tag also causes the controller to enable the circuitry in the first tag to transmit a number of symbols at an intersymbol period such that a number of clock cycles of the first tag during a transmission and receipt of the
25 symbols is substantially equal to a number of clock cycles of the clock of the second tag within an offset (n), where n and the number of symbols have no common factors other than one.

The first tag further includes a power source
30 coupled to the controller and supplying power to the circuitry in the first tag. The program in the memory of the first tag causes the controller to adjust the power supplied to the controller to that minimal level which allows communication between the circuitry in the first
35 tag and the circuitry in the second tag.

- 5 -

The system also includes a set of tags neighboring the first tag and including the second tag. A program in the memory of the first tag causes the controller to identify the set of tags by enabling the circuitry in the first tag to establish communications with each tag in the set. The controller determines the distance between each pair of tags in the set and the relative locations of each tag in the set based on the distances. The controller then selects a particular tag located on the periphery of the set for communicating outside the set.

In some embodiments, the information communicated between the tags includes a distance of the first tag from three different tags in the set and the absolute locations of the three different tags. The absolute locations may be derived from a sensor communicating its location to one of the three different tags. A program in the memory of the first tag causes the controller to determine the absolute position of the first tag from the information. The controller also minimizes an error in the absolute position of the first tag.

A program in the memory of the first tag also causes the controller to enable the circuitry in the first tag to communicate information with each tag in the set and analyze the information from the circuitry to determine the presence of a new tag in the set. The controller then causes the circuitry to communicate with the new tag to obtain a distance between the first tag and the new tag and a location of the new tag. The controller refines the absolute position of the first tag based on the location of the new tag and the distance between the first tag and the new tag.

In some embodiments, the information communicated by the tags includes the distance of the first tag from the second tag at a number of different times, and the absolute locations of the second tag at each of the

- 6 -

times. A program in the memory of the first tag causes the controller to determine the absolute position of the first tag from the information.

A program in the memory of the first tag causes
5 the controller to select tags in the set as relay tags, the circuitry in the relay tags being designated to route messages between a pair of tags in the set. The relay tags provide communications between each pair of tags in the set with a minimum number of communication links.
10 The controller performs a Delaunay triangulation to determine a topology of the tags in the set and selects the relay tags based on the topology.

In some embodiments, the system includes a user-operated device for communicating with the circuitry in
15 the first tag. The information communicated between tags includes an identity of a tag in the set capable of long range communication and a change in an identity of a tag in the set. The program in the memory of the first tag configures the controller to enable the circuitry in the
20 first tag to notify the device of the change in the identity of a tag in the set.

A program in the memory of the first tag also causes the controller to divide the set of tags into subsets when a number of tags in the set increases above
25 a threshold value. The memory in the first tag further stores an identification code for the first tag, and the program causes the controller to truncate the code and employ the circuitry in the first tag to communicate the truncated code to a second tag within a subset to
30 minimize energy consumption. The information communicated between tags may also include an identification of a tag belonging to a different set. A program in the memory of the first tag causes the controller to merge the first-mentioned set with the

- 7 -

different set when a number of tags in the first-mentioned set falls below a selected threshold value.

The system may also include a time sensor which communicates a time to the circuitry in the second tag.

5 The controller enables the circuitry in the first tag to communicate with the second tag to obtain the time, and sets the counter in the first tag using the communicated time.

A sensor may also communicate location information
10 to the first tag, and the memory in the second tag may store a transportation route for the object attached to the first tag. The controller employs the circuitry in the first tag to query the second tag regarding the transportation route and to query the sensor regarding
15 the tag's location. The controller then derives a present time from the counter in the first tag, determines an expected position of the object at the present time based on the route and an actual position of the object based on the location. The tag then decides
20 whether the actual position deviates from the expected position. The controller notifies the user-operated device of a deviation from the expected position by causing the circuitry to communicate with the device.

The user-operated device includes a processor
25 connected to a device communication circuitry. A device memory is coupled to the processor and includes a program configuring the processor to perform a called-for function when the processor reads the device memory. The function can cause the processor to determine a new route
30 for the object attached to the first tag based on the deviation, and causes the processor to employ the device communication circuitry to communicate the new route to the first tag. The processor also employs the device communication circuitry to receive communications

- 8 -

indicating a change in the route from at least one of the following: a second device, a tag and a system operator.

In some embodiments, the first tag has a detector connected to the controller. The detector notifies the controller when the first tag is detached from the object to which it had been attached. The controller enables the circuitry in the first tag to notify the user-operated device in response to notification from the detector that the first tag is detached from the object.

10 The circuitry in the first tag includes ports for infrared, radio frequency and acoustic communications. The port for radio frequency communications includes circuitry for communicating over a number of channels. A program in the memory causes the controller to select one channel for RF communication over the port and to code the information with an encryption code before enabling the tag's circuitry to communicate the information to the second tag.

The first tag can also include a power source connected to the controller and a timer connected to both the power source and controller. The controller loads a time period in the timer, and the timer causes the power from the power source to be removed from the controller and restored after the time period elapses.

25 A detector connected to the controller and the power source provides a detection signal to the controller indicating whether the power source has deteriorated beyond a threshold level. Another detector connected to the controller provides a detection signal to the controller indicating whether the tag has been tampered with. This detector includes a transmitter, a receiver, and an optical fiber connecting an output of the transmitter to an input of the receiver. The transmitter forwards a signal over the optical fiber to the input of the receiver. The receiver notifies the

- 9 -

controller if the signal does not arrive at the input of the receiver.

The system may also include a sensor detecting the condition of an object and communicating the condition to the circuitry in the first tag. The circuitry in the first tag subsequently communicates the condition to the second tag. The memory in the first tag includes a stored set of thresholds values related to the condition, each threshold being associated with a time. A program in the memory configures the controller to compare the condition to the threshold corresponding to a given time and to employ the circuitry to report a deviation from the threshold to the second tag. Similarly, the stored thresholds may be associated with a location of the first tag.

The sensor may be an analog sensor, such as a temperature sensor or a fuel level sensor for a vehicle. Alternatively, the sensor may be a digital sensor. If the object is disposed in a vehicle, the digital sensor may detect whether a door of the vehicle is open or closed, or whether a vehicle brake has been released.

In some embodiments, the system includes a computer having circuitry in communication with the circuitry in the first tag. A processor is connected to the circuitry in the computer, and a port for communicating with a user is coupled to the processor. A computer memory linked to the processor includes a stored threshold and a program configuring the processor to perform a selected function when the processor reads the computer memory. The function can cause the processor to enable the circuitry to receive information from the first tag, compare the information to the threshold, and enable the port to notify the user if the information exceeds the threshold.

- 10 -

The system can also include a module comprising a port for communicating with the circuitry in the first tag.

A processor is connected to the module port. A module
5 memory coupled to the processor includes a program
configuring the processor to perform a selected function
when the processor reads the module memory. The function
causes the processor to employ the port to receive an
activating signal and to provide information regarding an
10 object in the system in response to the activating
signal. The port includes circuitry for receiving a
wireless activating signal and circuitry for responding
to the activating signal with an acoustic signal. The
circuitry transmits the acoustic signal at a carrier
15 frequency identical to (or different from) a carrier
frequency of the activating signal.

In some embodiments, the system has a device with
an information exchange port configured to communicate
with the circuitry on the first tag. The port enables
20 the device to store information in the memory of the
first tag and to retrieve information from the memory of
the first tag. The device also includes circuitry for
long-range wireless communications with a wireless
component.

25 In general, in another aspect, the invention
features a module including a first tag having a port.
Circuitry in the first tag is coupled to the port and
configured for communicating with the second tag. A
second tag in the module includes a mechanism for
30 attaching the second tag to the first tag. A port
disposed on the second tag is adjacent the port on the
first tag when the second tag is attached to the first
tag. A circuit in the second tag is coupled to the port
in the second tag, the circuit being configured for
35 communicating with the circuitry in the first tag.

- 11 -

Preferred embodiments of this aspect of the invention include the following features.

The first tag includes a controller connected to the circuitry and a power connection coupled to the controller. The second tag includes a power source
5 coupled to its circuit and a second power connection. The second power connection couples to the first power connection when the second tag is attached to the first tag. The power source may include a connection to a
10 power source external to the second tag.

The second tag has an interrupt source coupled to a second port on the second tag. The first tag further includes an interrupt port coupled to its controller. The second port on the second tag is coupled to the
15 interrupt port on the first tag when the second tag is attached to the first tag.

The second tag includes a second circuit for communicating with a third tag and a switch connecting the second circuit to the first-mentioned circuit in the
20 second tag. The first tag includes a memory coupled to its controller which stores a program configuring the controller to execute a selected function when the controller reads the memory. The selected function causes the controller to communicate with the second tag
25 over the first-mentioned port and causes the second circuit in the second tag to communicate with the third tag. The memory may also cause the first tag's controller to monitor the second circuit in the second tag over the first-mentioned communication port for the
30 receipt of an activating signal. The controller causes the second circuit to respond with voice to the activating signal.

In general, in another aspect, the invention features a method for determining a topology of a set of
35 spatially distributed nodes. The method includes the

- 12 -

step of causing each node in the set to find its neighboring nodes by performing a Delaunay triangulation. Each node then stores an identification of each of its neighboring nodes.

5 A preferred embodiment of this aspect of the invention includes the following feature. A new node is added to the set and a Delaunay triangulation is performed at the new node to find the neighboring nodes of the new node. An identification of each of the
10 neighboring nodes is stored at the given node. Each neighboring node of the new node stores an identification of the new node.

 In general, in another aspect, the invention features a method for verifying a measured distance
15 between a pair of objects in a set of spatially distributed objects. The method includes the steps of choosing a group of three objects and measuring the distances between each pair of objects in the group. The method estimates an error for each distance and
20 determines whether the distances correspond to three sides of a triangle, taking into account the estimated errors.

 Preferred embodiments of this aspect of the invention include the following features.

25 The method further includes the steps of choosing a group of four objects from the set and repeating the measuring and estimating steps for each pair of objects in the group of four objects. The method then determines whether the distances correspond to the sides of a planar
30 tetrahedron sharing a common side, taking into account the estimated errors.

 The method may also include choosing a group of five objects from the set and repeating the measuring and estimating steps for each pair of objects in the group.
35 The method then determines whether the distances

- 13 -

correspond to the sides of a pair of planar tetrahedrons sharing a common triangle, taking into account the estimated errors.

In general, in another aspect, the invention
5 features a method of measuring a distance between two objects. The method includes the steps of transmitting from one of the objects a plurality of symbols, the transmissions occurring in a first series in which consecutive transmissions are separated by a first time
10 interval which is a multiple of a first clock period; receiving each of the symbols at the other of the objects and re-transmitting each of the symbols back to the one of the objects, the re-transmissions occurring in a second series in which consecutive re-transmissions are
15 separated by a second time interval which is different from the first time interval and which is a multiple of a second clock period which is different from the first clock period; measuring a round-trip time corresponding to detection of receipt from the other of the objects of
20 each of the symbols, each round-trip time being a multiple of the first clock period; and determining the distance between the objects using a velocity of the symbols, the clock periods, at least one of the time intervals, and the round-trip times.

25 Preferred embodiments of this aspect of the invention include the following features. One clock period may be approximately equal to the product of the other clock period and a fraction, the numerator of the fraction may be equal to the sum of two terms, the first
30 term of the sum may be the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval, the second term of the sum may be a whole number that shares no multiplicative factors, other than 1, with the number of
35 symbols transmitted, and the denominator of the fraction

- 14 -

may be equal to the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval. The measured round-trip times may vary and may have a variation range
5 equal to one first clock period.

In another aspect, the invention features another method of measuring a distance between two objects. The method includes the steps of transmitting from one of the objects a plurality of symbols, the transmissions
10 occurring in a series in which consecutive transmissions are separated by a time interval which is a multiple of a first clock period; returning each of the symbols from the other of the objects back to the one of the objects; measuring a round-trip time corresponding to detection of
15 receipt from the other of the objects of each of the symbols, each round-trip time corresponding to a multiple of a second clock period; and determining the distance between the objects using a velocity of the symbols, the clock periods, at least one of the time intervals, and
20 the round-trip times.

Preferred embodiments of this aspect of the invention include the following features. The method may further include a step of detecting each transmission using the second clock period, wherein the round-trip
25 times correspond to time intervals between detections of receipt and transmission of the symbols. One of the clock periods may be approximately equal to the product of the other clock period and a fraction, the numerator of the fraction may be equal to the sum of two terms, the
30 first term of the sum may be the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval, the second term of the sum may be a whole number that shares no multiplicative factors, other than 1, with the number of
35 symbols transmitted, and the denominator of the fraction

- 15 -

may be equal to the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval. The measured round-trip times may vary and may have a variation range equal to one first clock period.

In another aspect, the invention features another method of measuring a distance between two objects. The method includes the steps of transmitting from one of the objects a first symbol; transmitting from the one of the objects a second symbol, the transmission of the second symbol being separated from the transmission of the first symbol by a time interval which is a multiple of a first clock period; receiving the first symbol at the other of the objects and re-transmitting the first symbol back to the one of the objects; receiving the second symbol at the other of the objects and re-transmitting the second symbol back to the one of the objects, the re-transmission of the second symbol being separated from the re-transmission of the first symbol by a second time interval which is a multiple of a second different clock period; measuring a first round-trip time corresponding to detection of receipt from the other of the objects of the first symbol, the first round-trip time being a multiple of the first clock period; measuring a second round-trip time corresponding to detection of receipt from the other of the objects of the second symbol, the second round-trip time being a multiple of the first clock period; if the round-trip times are not different by at least one first clock period, repeating the transmitting, the receiving, the re-transmitting, and the round-trip measuring of the second symbol until the round-trip times are different by one first clock period; and determining the distance between the objects using a velocity of the symbols, the clock periods, the time intervals, and the round-trip times.

- 16 -

A preferred embodiment of this aspect of the invention includes the following feature. The first and second round-trip times may be original round-trip times, the original round-trip times being different by one
5 first clock period, and the method may further include the steps of (a) equating the clock periods; (b) temporarily changing one of the clock periods for a period of time equal to a fraction of the time interval; (c) determining a new round-trip time for a symbol; (d)
10 comparing the new round-trip time to at least one of the original round-trip times; and repeating (a) through (d) for a plurality of iterations, the fraction being different for each iteration.

In another aspect, the invention features another
15 method of measuring a distance between two objects where one of the objects is able to measure time with a resolution of a minimum period of time. The method includes the steps of transmitting a plurality of symbols from the one of the objects; receiving each of the
20 symbols at the other of the objects and re-transmitting each of the symbols back to the one of the objects after a delay period, the delay period being different for each symbol and being a multiple of a constant delay unit, the delay unit being approximately equal to the minimum
25 period divided by the number of symbols transmitted; measuring a round-trip time corresponding to detection of receipt of each symbol from the other of the objects, each round-trip time corresponding to a multiple of the minimum period; and determining the distance between the
30 objects using a velocity of the symbols, the minimum period, and at least one of the round-trip times.

In another aspect, the invention features another method of measuring a distance between two objects, where one of the objects is able to measure time with a
35 resolution of a minimum period of time. The method

- 17 -

includes the steps of transmitting a plurality of symbols from the one of the objects; after a first delay period, detecting each transmission, the first delay period being different for each transmission and being a multiple of a
5 constant delay unit, the delay unit being approximately equal to the minimum period divided by the number of symbols transmitted; returning each of the symbols at the other of the objects back to the one of the objects; after a second delay period, detecting the receipt of
10 each of the symbols from the other of the objects, the second delay period being different for each receipt and being a multiple of the delay unit; measuring a time difference between detections of transmission and receipt of each of the symbols, each time difference being a
15 multiple of the minimum period; and determining the distance between the objects using a velocity of the symbols, the minimum period, and the time differences.

In another aspect, the invention features another method of measuring a distance between two objects, where
20 one of the objects is able to measure time with a resolution of a minimum period of time. The method includes the steps of transmitting a first symbol from the one of the objects; receiving the first symbol at the other of the objects and re-transmitting the first symbol
25 back to the one of the objects; measuring a first round-trip time corresponding to detection of receipt from the other of the objects of the first symbol, the first round-trip time being a multiple of the minimum period;
a) transmitting another symbol from the one of the
30 objects; b) receiving the another symbol at the other of the objects and re-transmitting the another symbol back to the one of the objects after a delay period, the delay period being a multiple of a delay unit, the delay unit being approximately equal to a fraction of the minimum
35 period, the number of delay units in the delay period

- 18 -

being equal to the number of symbols previously re-transmitted; c) measuring another round-trip time corresponding to detection of receipt from the other of the objects of the another symbol, the another round-trip time being a multiple of the minimum period; if the first
5 round-trip time and the another round-trip time are equal, repeating a) through c) until the round-trip times are different by one minimum period; and determining the distance between the objects by using a velocity of the
10 symbols, the minimum period, delay unit, and the round-trip times.

A preferred embodiment of this aspect of the invention includes the following feature. The method may further include a step of reducing the delay unit to a
15 smaller fraction of the minimum period and repeating a) through c).

In general, in another aspect, the invention features apparatus for measuring a distance between two objects. The apparatus includes a transmitter having a
20 first clock for transmitting from one of the objects a plurality of symbols, the transmissions occurring in a first series in which consecutive transmissions are separated by a first time interval which is a multiple of the period of the first clock; a transponder having a
25 second clock for receiving each of the symbols at the other of the objects and re-transmitting each of the symbols back to the one of the objects, the re-transmissions occurring in a second series in which consecutive re-transmissions are separated by a second
30 time interval which is different from the first time interval and which is a multiple of the second clock's period which is different from the period of the first clock; a clock counter for measuring each round-trip time corresponding to detection of receipt from the other of
35 the objects of each of the symbols, each round-trip time

- 19 -

being a multiple of the first clock period; and a calculator for determining the distance between the objects using a velocity of the symbols, the clock periods, at least one of the time intervals, and the
5 round-trip times.

In another aspect, the invention features another apparatus for measuring a distance between two objects. The apparatus includes a transmitter having a first clock for transmitting from one of the objects a plurality of
10 symbols, the transmissions occurring in a series in which consecutive transmissions are separated by a time interval which is a multiple of the period of the first clock; a transponder for returning each of the symbols from the other of the objects back to the one of the
15 objects; a counter having a second clock for measuring a round-trip time corresponding to detection of receipt from the other of the objects of each of the symbols, each round-trip time corresponding to a multiple of the period of the second clock; and a calculator for
20 determining the distance between the objects using a velocity of the symbols, the clock periods, the time interval, and the round-trip times.

In another aspect, the invention features another apparatus for measuring a distance between two objects.
25 The apparatus includes a transmitter for transmitting from one of the objects a first symbol and transmitting from the one of the objects a second symbol, the transmission of the second symbol being separated from the transmission of the first symbol by a time interval
30 which is a multiple of a first clock period; a transponder for (a) receiving the first symbol at the other of the objects and re-transmitting the first symbol back to the one of the objects and (b) receiving the second symbol at the other of the objects and re-
35 transmitting the second symbol back to the one of the

- 20 -

objects, the re-transmission of the second symbol being separated from the re-transmission of the first symbol by a second time interval which is a multiple of a second different clock period; a clock counter for (a) measuring
5 a first round-trip time corresponding to detection of receipt from the other of the objects of the first symbol, the first round-trip time being a multiple of the first clock period, and (b) measuring a second round-trip time corresponding to detection of receipt from the other
10 of the objects of the second symbol, the second round-trip time being a multiple of the first clock period; a comparator for (a) testing if the round-trip times are not different by one first clock period, and (b) initiating repeating of the transmitting, the receiving,
15 the re-transmitting, and the round-trip measuring of the second symbol until the round-trips are different by one first clock period; and a calculator for determining the distance between the objects using a velocity of the symbols, the clock periods, the time intervals, and the
20 round-trip times.

In general, in another aspect, the invention features another method of measuring a distance between two objects. The method includes the steps of transmitting from one of the objects a plurality of
25 symbols, the transmissions occurring in a first series in which consecutive transmissions are separated by a first time interval which is a multiple of a first clock period, the symbols having a constant velocity; receiving each of the symbols at the other of the objects and re-
30 transmitting each of the symbols back to the one of the objects, the re-transmissions occurring in a second series in which consecutive re-transmissions are separated by a second time interval, the second time interval being different from the first time interval,
35 the second time interval being a multiple of a second

- 21 -

clock period which is different from the first clock period, each re-transmission thus occurring after a delay, the delay being different for each re-transmission and constituting the period of time between arrival of
5 each symbol at the other of the objects and re-transmission of each symbol, the delay having a nominal length which is a multiple of a delay unit, the delay unit being equal to the product of the number of first clock periods corresponding to the first time interval
10 and the difference between the clock periods, the delay unit also being equal to the second clock period divided by the number of symbols in the plurality of symbols, the delay having an actual length which is within one delay unit of the nominal length; measuring a round-trip time
15 corresponding to detection of receipt from the other of the objects of each of the symbols, each round-trip time being a multiple of the first clock period; and determining the distance between the objects by multiplying the velocity of the symbols by half of the
20 difference between the average of all of the round-trip times and the first clock period.

Preferred embodiments of this aspect of the invention include the following features. The delay may further include an extra delay which is a multiple of the
25 second clock period; and the distance may be determined by multiplying the velocity by half of the difference between (a) the average of all of the round-trip times and (b) the sum of the extra delay and the first clock period. The method may further include a step of
30 offsetting the average of the round-trip times, wherein the offset is an amount of time equal to the second clock period divided by twice the product of (a) the number of symbols and (b) the number of first clock periods corresponding to the first time interval.

- 22 -

In another aspect, the invention features another method of measuring a distance between two objects. The method includes the steps of transmitting from one of the objects a plurality of symbols, the transmissions
5 occurring in a series in which consecutive transmissions are separated by a time interval which is a multiple of a first clock period, the symbols having a constant velocity; detecting the transmissions using a second clock period, the second clock period being different
10 from the first clock period, each transmission detection thus occurring after a delay, the delay being different for each detection and constituting the period of time between each transmission and the detection of each transmission, the delay having a nominal length which is
15 a multiple of a delay unit, the delay unit being equal to the product of the number of first clock periods corresponding to the time interval and the difference between the clock periods, the delay unit also being equal to the second clock period divided by the number of
20 symbols in the plurality of symbols, the delay having an actual length which is within one delay unit of the nominal length; returning each of the symbols from the other of the objects back to the one of the objects; detecting the arrivals of the returning symbols back at
25 the one of the objects using the second clock period, each arrival detection thus occurring after a second delay, the second delay being different for each detection and constituting the period of time between each arrival and the detection of each arrival, the
30 nominal length of the second delay being a multiple of the delay unit, the second delay having an actual length which is within one delay unit of the nominal length of the second delay; measuring a round-trip time corresponding to detection of receipt from the other of
35 the objects of each of the symbols, each round-trip time

- 23 -

corresponding to a multiple of a second clock period; and determining the distance between the objects by multiplying the velocity by half of the average of the round-trip times.

5 A preferred embodiment of this aspect of the invention includes the following feature. The symbols may be returned from the other of the objects after an extra delay; and the distance may be determined by multiplying the velocity by half of the difference
10 between the average of all of the round-trip times and the extra delay.

Advantages of the invention include the following.

An operator is required to enter data into the system only once. After entry, the tags and devices in
15 the system automatically move the data to where it is most needed. For example, the operator enters the contents of a particular asset into the tag attached to the asset. The tag then automatically communicates the contents to an inventory computer storing a list of local
20 assets.

Because the tags communicate with other tags in their vicinity, the effective communications range of each tag extends beyond the range of its own communication circuitry. The effective range is further
25 extended by placing an additional tag near a selected tag in the system. The selected tag relays communications from the remaining tags in the system to the additional tag. The remaining tags are thus not required to be located within the communication range and line of sight
30 of the additional tag. Once the tags are activated, they automatically seek other tags or devices in order to report their status. The system thus does not need to depend on the honesty of company employees in collecting and reporting data relating to the assets. The
35 information reported by the tags takes a variety of

- 24 -

forms. For example, a tag may monitor a sensor for the temperature of an object, and may notify a device of an abnormal increase in temperature. Similarly, the tag may determine a scheduled route for the asset, and report any
5 deviations from that route to the user-operated device.

An operator obtains information relating to an asset in the system by directly querying an appropriate tag. Because several tags in the system share the information and can automatically move the information,
10 the operator is likely to have access to a tag that either stores the needed information or that can acquire the information. The operator is thus not required to access an operations center that may be outside the operator's communications range in order to obtain the
15 information.

A local or remote computer is easily integrated into the system. The computer gathers and disseminates information to operators who do not possess the appropriate device for querying tags, or to operators who
20 are located far from any tag. The computer maintains a list of tags, their location and associated asset, so that any operator query can be directed to the tag.

Unlike the barrier systems described above, the system provides continuous spatial monitoring of the
25 assets. Instead of only determining that a particular asset is located anywhere inside one or more cells, the system accurately locates a position of the asset within any cell. In addition, while barrier systems cannot typically determine a direction of movement of an asset
30 when it crosses a boundary, the system according to the invention accurately determines direction and velocity of the asset. Knowing the position and velocity of the asset allows operators to promptly find the asset in response to an emergency, and permits accurate policing
35 of the asset to prevent future incidents. The assets are

- 25 -

thus more secure from theft or misplacement, which is an important consideration when the assets include valuable cargo. In addition, the quick containment of incidents provides for more safety in transporting and storing of the assets. This is particularly important for assets carrying hazardous materials.

In some embodiments, the tags form an optimal communications network without user input, and update the network whenever there is a change in the location or identity of the assets. The tags define the network by designating tags for relaying information between specific pairs of tags. The network is optimized by minimizing the number of relay tags between every pair while maximizing the range of the communications. The tags superimpose on this network a hierarchy of communication links to other networks or devices.

The tags also optimize the network by assigning specific functions to particular tags in order to prevent tags from needlessly repeating the functions of other tags. These functions include checking whether assets have been added to or removed from the network.

The tags design the network based on their calculations of the location of other tags. Error correcting algorithms in the tags permit them to compute these locations with a high degree of accuracy. When new information becomes available to a tag (for example, when a new tag is added to the network), the tags automatically refine their location calculations.

When a network becomes too large, the tags automatically optimize the network by dividing it into subnetworks. When a network becomes too small, the tags join the network to another network. The communications links include existing wireless communication links, for example, cellular telephone links. The system thus does not require investment in a

- 26 -

long distance communications infrastructure. The system is also flexible enough to use a number of communication protocols. This is accomplished by providing tags to translate communications between different standards, such as are used with radio frequency, infrared, sound and serial link.

The tags remain active for long periods by employing a variety of energy conservation methods. These include tapping an external power source (for example, a vehicle battery) directly or through an intermediary tag, or by including a large or rechargeable battery in the tag. The tags power themselves down between operations. This is possible because the tags can process and communicate information faster than they are typically queried by an operator. In a powered down, comatose mode, the batteries consume energy at a rate comparable to the degeneration rate of the stored battery.

The tags communicate with low power by minimizing the time required for communications, and using the optimal relay tags designated in the network. The tags further conserve energy by only reporting the existence of exceptional conditions to the device, instead of continuously reporting normal conditions. For example, a tag only communicates the fact that a temperature of an asset exceeds a given threshold, instead of continually reporting the temperature to the device. Once subnetworks are formed, the tags conserve energy by communicating compacted information within the subnetwork, and uncompacting the information for communication outside the network. For example, each tag in the subnetwork is assigned a truncated ID that is shorter than its actual, global ID. The code indicating a type of communications between the tags is similarly truncated. The tags within a subnetwork synchronize

- 27 -

their functions by powering on for a short time simultaneously. Arranging the tags in a hierarchy of networks requires that only a selected tag within a network perform long range communications with a tag in another network. The tags within a network, being
5 located within closer range of one another, employ lower energy for all their communications. The selected tag communicates only with tags in a network at an adjacent level of the hierarchy. The tag in the adjacent level
10 acts as a repeater, relaying the selected tag's message to another level of the hierarchy. The selected tag thus does not perform the longer range communications necessary to communicate directly with a network removed by several levels.

15 The methods and apparatus used for finding the distance between two objects or tags are efficient in time and in the number of steps involved. Significant improvements in distance measurement are obtained through the use of only a handful of symbols.

20 For distance measurement, the two objects operate in half-duplex mode so that separate channels are not needed for the outgoing leg and the return leg of the symbols' round-trip.

The distance measurement is made using only two
25 clocks, which in some embodiments are included in the same object. If the clocks are included in the same object, the other object need operate only as a simple reflector or repeater of the symbols.

The methods and apparatus used for finding the
30 distance between two objects also consume little energy.

The tags are standardized to allow them to be easily added and removed from a network. Flexibility is provided by connecting the tags to special tags (helper tags), which perform specialized support functions.
35 These functions include providing additional power to the

- 28 -

attached tag and acting as a communications interface for the tag.

The tags are small and easily fixed to assets. The lifetime and cost (including labor cost of use) of the tags is comparable to that of bar codes. The user interfaces on the user-operated monitoring device and computers are simple to use and reduce operator errors. For example, the monitoring device includes a personal digital assistant which has a menu of options. Enabling the user to select options instead of entering commands reduces the user's workload and minimizes potential errors.

Other features and advantages of the invention will become apparent from the following description and claims.

Brief Description of the Drawings

Fig. 1 is a schematic representation of an asset management system according to the invention.

Fig. 2 is a schematic representation of a collection of assets in the system of Fig. 1.

Fig. 3 is a schematic representation of the system of Fig. 1 as applied to a warehouse storing assets.

Fig. 4 is a schematic representation of the system of Fig. 1 as applied to a warehouse storing hazardous materials.

Fig. 5 is a schematic representation of the system of Fig. 1 as applied to a train carrying assets.

Fig. 6 is a schematic representation of the system of Fig. 1 as applied to a truck carrying assets.

Fig. 7 is a schematic representation of the system of Fig. 1 as applied to a shipping port moving assets between a ship and a holding area.

Fig. 8 is a schematic representation of the system of Fig. 1 as applied to a bonded carrier moving assets.

- 29 -

Fig. 9 is a schematic representation of the system of Fig. 1 as applied to a clandestine tracking system for stolen or contraband assets.

Fig. 10 is a perspective view of an asset tag and helper tag in the system of Fig. 1.

Fig. 11 is a block diagram of an asset tag port on the helper tag of Fig. 10.

Fig. 12 is a block diagram of communications and power control circuitry in the asset tag of Fig. 10.

Fig. 13 is a block diagram of detector, memory and power control circuitry in the asset tag of Fig. 10.

Fig. 14 is a block diagram of a sensor in the asset tag of Fig. 10.

Fig. 15 is a block diagram of a helper tag in the system of Fig. 2.

Figs. 16 and 17 are block diagrams of an interrogate and respond module in the system of Fig. 2.

Fig. 18 is a schematic representation of a monitoring device in the system of Fig. 1.

Fig. 19 is a block diagram of a helper tag in the device of Fig. 18.

Fig. 20 is a block diagram of a holster in the device of Fig. 19.

Fig. 21 is a block diagram of a communications controller in the device of Fig. 19.

Fig. 22 is a block diagram of software in the controller of Fig. 21.

Figs. 23 through 25 are block diagrams of functions performed by a processor in the asset tag of Fig. 1.

Figs. 26 through 28 are flowcharts for software performed by a self-assessment unit in the processor of Fig. 23.

Figs. 29 and 30 are flowcharts for software performed by a joining unit in the processor of Fig. 23.

- 30 -

Fig. 31 is a flowchart for an software performed by a location unit in the processor of Fig. 23.

Figs. 32 and 33 are flowcharts for software performed by an agent unit in the processor of Fig. 23.

5 Fig. 34 is a schematic diagram of two tags in the system of Fig. 1.

Fig. 35 is a graph of signals in the asset tags of Fig. 34.

10 Fig. 36 is a second graph of signals in the asset tags of Fig. 34.

Fig. 37 is a flowchart for a process performed by the tags in Fig. 34.

Figs. 38 and 39 are graphs of errors in a range measured by the asset tags of Fig. 34.

15 Fig. 40 is a flowchart for a process performed by the tags in Fig. 34.

Fig. 41 is a graph of measurement results in the tags of Fig. 34.

20 Fig. 42 is a flowchart for a process performed by the tags in Fig. 34.

Fig. 43 is a graph of signals in the tags of Fig. 34.

Fig. 44 is a schematic representation of a list stored in a memory of the asset tag of Fig. 1.

25 Fig. 45 is a flowchart for software performed by the processor of the asset tag in Fig. 1.

Fig. 46 is a schematic representation of group of three asset tags in the system of Fig. 1.

30 Fig. 47 is a schematic representation of a list stored by an asset tag in Fig. 1.

Fig. 48 is a schematic representation of 4 tags in the system of Fig. 1.

Figs. 49 and 50 are schematic representations of lists stored by an asset tag in Fig. 1.

- 31 -

Fig. 51 is a schematic representation of a group of tags in the system of Fig. 1.

Fig. 52 is a flowchart for software performed by an asset tag in Fig. 1.

5 Figs. 53 and 54 are schematic representations of a group of tags in the system of Fig. 1.

Figs. 55 and 56 are flowcharts for software performed by an asset tag in Fig. 1.

10 Figs. 57 and 58 are schematic representations of a group of tags in the system of Fig. 1.

Figs. 59 and 60 are flowcharts for software performed by an asset tag in Fig. 1.

Fig. 61 is a schematic representations of a group of tags in the system of Fig. 1.

15 Fig. 62 is a flowchart for software performed by an asset tag in Fig. 1.

Figs. 63 and 64 are schematic representations of the application of Fig. 6.

20 Fig. 65 is a schematic representation of functions performed by the asset tags in Fig. 63.

Fig. 66 is a schematic representation of the railroad application of Fig. 5.

Figs. 67 and 68 are schematic representations of groups of tags in the system of Fig. 1.

25 Fig. 69 is a flowchart for software performed by an asset tag of Fig. 1.

Description of the Preferred Embodiments

I. Overview of the Asset Management System

Referring to Fig. 1, an asset management system 10
30 includes a collection of assets 12, individual assets being (identified as $12_1, 12_2, \dots 12_n$) each having an asset tag 16 (identified as $16_1, 16_2, \dots 16_n$) linked to a remote operations center 13 and a local computer 14. The operations center and computer both store and process
35 information relating to the assets. A remote user

- 32 -

employs the operations center to access the information, while a local user queries the computer. A local monitoring device 15 is linked to both the collection of assets and the operations center, and provides a mobile, portable interface for the local user.

The operations center, local computer and monitoring device can be linked to the collection of assets by at least one communication relay, for example, a cellular base station. Alternatively, the center, computer and device are linked directly to the collection of assets.

Referring to Fig. 2, in one illustrated embodiment, the assets in the collection includes items $12_2, 12_3, \dots, 12_5$, a container 12_1 storing items $12_2, 12_3, 12_4$, and a vehicle 12_6 for transporting the container and the items. For example, in a monitoring system for a warehouse, the assets include forklifts transporting pallets that store containers.

Each asset in the collection is equipped with an active electronic asset tag $16_1, 16_2, \dots, 16_6$. Each asset tag is in communication with other asset tags that are within their communication range. A partial set of possible communication links 17 between asset tags is diagrammatically shown in Fig. 2 although in actuality, the links would be wireless. Additional asset tags $16_7, 16_8$ can be placed in the vicinity of the assets to assist in tag communications.

Selected asset tags can also communicate with sensors $18_1, 18_2$ and helper tags $20_1, 20_2, 20_3$ located near the assets. The sensors monitor a characteristic of an asset or of the asset's environment. For example, sensor 18_1 monitors the temperature of item 12_5 , while sensor 18_2 detects the location of the vehicle 12_6 . The helper tags perform a variety of specialized functions. For example, helper tag 20_1 provides power to asset tag 16_7 , and helper

- 33 -

tag 20₂ acts as a communication interface for asset tag 16₆. Asset tag 16₈ and helper tag 20₁ together form an interrogate and respond module 22 that provides a description of the collection of assets when activated by an interrogator, for example, a radar gun. Some asset tags communicate with the monitoring device 15 and the computer 14 located near the assets (Fig. 1). The monitoring device and some of the tags and sensors are additionally capable of long range communication with the remote operations center 13.

Some or all of the asset tags and helper tags process information regarding other asset tags and sensors in their vicinity, and use this information to form a communication network. The network defines a set of communication links between a neighboring group of helper tags, asset tags, sensors, devices and computers that maximizes the efficiency of communications within the group.

The asset tags are either battery powered or powered by an external source. The battery powered tags include various mechanisms for minimizing the energy consumed by the tag. For example, the tags are placed in a low energy, sleep state between communications.

The local user employs the monitoring device 15 to survey the status of the assets in a variety of ways (see Fig. 1). For example, the local user may attempt to communicate with asset tag 16₅ to determine whether the attached item 12₅ is present in the system. In addition, the local user may use the device to query sensor 18₂ through helper tag 20₂ to determine a location of vehicle 12₆.

The remote user may inform the operations center 13 of a new route planned for the transportation of a particular item in the system. The center relays this information to the monitoring device, which downloads the

- 34 -

information to the tag attached to the particular item. The local user learns of the new route through the monitoring device.

The general system described above is tailored to specific applications by omitting certain components of the system, and/or by modifying the functions performed by a particular component in the system. For example, an application may omit the operations center and rely only on the local computer. The monitoring device in this system would then omit all functions relating to an operations center. The following section provides a short overview of several specific applications.

II. Overview of Specific Applications of the System

Referring to Fig. 3, one application 10₁ of the system 10 is designed for monitoring the location of containers 12₁ in a warehouse. (The assets will be uniformly labelled in the examples which follow as 12_i and the tags as 16_i.) Each container 12₁ in the warehouse, along with a forklift truck 12₆ used for moving the containers, is equipped with an asset tag 16₁, 16₆, respectively. Tag 16₆ on the truck is equipped with a global positioning system (GPS) sensor 18₂ for determining the truck's location.

An asset tag 16₉ (a gateway tag) capable of long range communications is mounted on a wall of the warehouse, and is in communication with a central inventory computer 14. The tags automatically form an efficient communication network for sending information to the inventory computer through the gateway tag. A user 50 monitors the location of the containers by using monitoring device 15 to directly access the asset tags on the containers. The user may also access the tags through the inventory computer.

- 35 -

In another application 10₂, shown in Fig. 4, a warehouse stores containers 12₁ of hazardous waste. Each container is equipped with an asset tag 16₁. Inventory computer 14 communicates with the asset tags through gateway tag 16_g, in a manner identical to that described above in connection with Fig. 3.

An interrogate and respond module 22 is attached to an outside wall 13 of the warehouse. In the event of an emergency, a fireman 52 informs himself of the contents of the warehouse by querying the module using a standard police radar gun 54. The module responds with a synthesized voice message over the fireman's radio receiver 55. The fireman is thus notified of the hazardous contents of the warehouse before entering the warehouse.

Referring to Fig. 5, another application 10₃ monitors railway cargo by mounting a tag 16_g on each railway car 12_g and on the locomotive 12_g'. Tags 16₁ are also mounted on each container 12₁ transported in the cars. The tags on the containers form a network with the tags on the cars and on the locomotive. The conductor uses a monitoring device 15 to communicate with the operations center 13 through the tag on the locomotive.

Several sensors 18 can be mounted, according to this embodiment, in the railway cars. These include a brake release sensor monitoring whether the brakes in a car have been released, a temperature sensor, and a fuel level sensor. A GPS sensor is typically provided near the tag on the locomotive.

A similar application 10₄, shown in Fig. 6, involves monitoring the transportation of hazardous waste or other materials by a truck 12_g. Asset tags 16_g, 16₁ are fixed to the truck and each container 12₁ in the truck. A driver communicates with the tags on the containers with a monitoring device 15, which reports

- 36 -

relevant information to an operations center 13. Unlike the railroad application described above, the tags in this application need not automatically form a communications network.

5 Another application 10₅, shown in Fig. 7, monitors the location of containers at a shipping port 60. Each container is equipped with an asset tag. Trucks 12₆ and trains 12₆' carrying the containers to and from a ship 62 and holding area 64, are equipped with the systems in
10 applications 10₃, 10₄ described above, respectively. The holding area 64 and the ship 62 include the warehouse application 10₁ described in connection with Fig. 3. Arrows in Fig. 7 show the flow of containers between the ship and holding area.

15 Referring to Fig. 8, another application 10₆ of the system monitors bonded carriers that transport goods near or across national borders. These carriers are required to prevent their conveyances from being opened as they travel between specified locations near the
20 borders.

Each container 12₁ in the conveyance is equipped with an asset tag 16₁ incorporating a tamper detector. When the tamper detector indicates to the asset tag that it has been damaged, the tag alerts an operations center
25 13. The operations center in turn notifies the carrier that one of its containers may have been tampered with.

Another application 10₇, shown in Fig. 9, monitors the location of valuable assets 12₁, 12₂ in a clandestine fashion to prevent theft. For example, the assets may
30 include jewelry in a home or art objects in a museum.

The tags 16₁', 16₂' on the assets are hidden from view, to prevent a thief from noticing and removing the tags. This is generally accomplished by attaching very small but relatively low power tags to the assets.
35 The tags are tracked by a monitoring device 15, or by

- 37 -

another tag 16₆, for example, a tag mounted in a conveyance in which the assets are being transported. The tag 16₆ notifies the monitoring device 15 or an operations center 13 of the presence of the asset in the
5 conveyance.

The full range of functions of each component in the asset management system of Figs. 1 and 2 is described in the following sections of the Description. When used in one of the applications shown in Figs. 3 through 9, a
10 component may omit some of these functions. For example, the asset tags in the application of Fig. 6 do not automatically network with other tags. Details of the particular design of the components and their interactions in particular applications are given in
15 Section XII.

III. Asset Tag

A. Structure of the tag

In the embodiment of the invention illustrated in Fig. 10, each tag 16 is a small rectangular package 2.5
20 cm on a side and 1 cm in thickness. (Other shapes and sizes could be used.) The illustrated tag has three recesses 96, 97, 98 on one face 99.

Referring to Fig. 11, the interior of each tag includes a microprocessor 100, for example, model
25 MC68L11A1FN available from *Wyle Electric, 20 Chapin Rd., Bldg. 1013, P.O. Box 684, Pine Brook, NJ 07058*, which is provided with an adjustable clock. The microprocessor is connected to communications circuitry 102 which performs short range infrared (IR) and long range radio frequency
30 (RF) communications with other tags and devices. The circuitry in this embodiment is connected to a helper tag port 103 located on face 99 of the tag (Fig. 10).

The microprocessor is coupled to a memory 104 for storing information and a set of detectors 106 for
35 determining a status of the tag. The microprocessor is

- 38 -

also coupled to an interrupt port 107 for receiving an interrupt signal from outside of the tag. The interrupt port 107 is connected at recess 98 on face 99 of the tag (Fig. 10).

5 A power controller 109 in the tag links the microprocessor to a power source 110. A timer 112 connected between the microprocessor and the power controller regulates the controller. The controller supplies power to the detectors, the memory, the timer
10 and the communication circuitry, adjusting the power supplied to the communications circuitry by up to 10 db as commanded by the microprocessor.

 Referring to Fig. 12, the communication circuitry 102 (shown here in greater detail) includes an RF
15 receiver 130 and transmitter 132 and an IR receiver 134 and transmitter 136, each of which is connected to the microprocessor 100.

 The microprocessor causes the RF receiver and transmitter to communicate at, for example, 916.5 MHz
20 with one of ten different spread spectrum, direct sequence gold codes. The IR transmitter and receiver communicate, preferably, with an unmodulated IR signal. The microprocessor also causes selected communications to be coded with public or private key encryption. The
25 maximum communication range of the RF circuitry is on the order of 100 m, at a data rate, typically, of 25 kbps. The power source 110 includes a battery 138 and an external power port 140. The battery has a two year active lifetime, with a ten year storage life. Power port 140
30 includes a power connection and a ground connection coupled to recesses 96, 97, respectively, on face 99 of the tag (Fig. 10).

 The power controller 109 has three separate controls: a device power control 142, a processor power
35 control 144 and a communications power control 148. The

- 39 -

battery and external power source are coupled to the device power control by a switch 150 which automatically switches to the external power port when an external power source is available.

5 The processor power control 144 is connected to the device power control 142 and provides power directly to the microprocessor. The processor power control in turn receives control signals from the microprocessor. The communications power control 148 also derives power
10 from the processor power control 144, and accepts control signals from the microprocessor. The communications power control 148 in turn supplies power to the IR transmitter 136, IR receiver 134, RF transmitter 132 and RF receiver 130.

15 Referring to Fig. 13, the processor power control 144 is additionally connected to the timer 112 by the device power control 142. The timer is in direct communication with the microprocessor.

 The memory 104 in the tag (Fig. 11) includes a
20 random access memory 160 (RAM) of at least 32 kB capacity, together with an identity store 162 (Fig. 13). Both the RAM and identity store are connected to the microprocessor, and receive power from the processor power control.

25 The detectors 106 in the tag (Fig. 11) include a detach detector 164 (Fig. 13) and a battery sensor 166. Both the detach detector and the battery sensor are connected to the microprocessor. The battery sensor is connected to the battery 138, and provides a signal to
30 the microprocessor when the battery voltage falls below 2.7 V in the illustrated embodiment. The detach detector includes an electromechanical switch placed on one face of the tag, that is depressed when the tag is mounted on an asset. The switch derives power from processor power
35 control 144 and provides a signal when the tag is

- 40 -

detached from the asset, causing the switch to be released. B. Basic operation of the tag

In use, the tag conserves energy by remaining in a low power state for approximately 99% of its lifetime.

5 The tag remains in one type of low power state (a comatose state) until activated by an interrupt signal over interrupt port 107 (Fig. 11) or by a low frequency signal from timer 112. A second type of low power state (a sleep state) is interrupted by a higher frequency
10 signal from the timer 112. The tag begins its functions when an operator couples the external power port on the tag to an external power source. This causes switch 150 to connect the port to the device power control.

15 When connected to an external power source, the processor power controller supplies power from the power port to the microprocessor, placing the microprocessor in an awake state. The microprocessor remains in the awake state until it receives a high voltage signal from the
20 interrupt port. The interrupt signal is typically supplied by connecting a helper tag 20 (Fig. 1) to the interrupt port, as described below. The microprocessor then performs various functions, including communicating with the monitoring device 15 (Fig. 1) or
25 other asset tags. These functions are described in detail in the following sections.

After the microprocessor has completed its functions, the microprocessor sets a frequency at which the timer sends a wake up signal to the power controller.
30 The timer notifies the power controller to reduce the power supplied, to place the microprocessor in a sleep state. The microprocessor is subsequently awakened from the sleep state each time the power controller receives a wake-up signal from the timer. The microprocessor
35 synchronizes the wake-up timer with other asset tags and

- 41 -

devices with which the tag communicates, as follows. First, the tag seeks an absolute time signal from a GPS sensor 18₂ (Fig. 2), to set the internal clock in the microprocessor to the absolute time. The internal clock
5 in the microprocessor records the length of time over which the tag remains awake, that is, the processing time required for the tag to communicate with other tags and devices, and to perform various functions. After being awakened several times, the microprocessor derives an
10 average processing time for each wake-up signal. Based on the average processing time, the microprocessor informs the timer of an ideal wakeup frequency. Subsequently, the timer supplies signals to the power controller causing the power controller to wake the
15 microprocessor from its sleep state at the ideal frequency.

When employing the RF transmitter, the microprocessor causes the transmitter to communicate an appropriate synchronization code for spread spectrum
20 communication. If the tag is not synchronized with other tags, the transmitter broadcasts the code for T seconds, allowing enough time for the other asset tags to wake up and receive the code.

C. Other Embodiments

25 In other embodiments, the communication circuitry in the asset tag includes acoustic or optical communications, or long range communications circuitry. The RF communication protocol used may be 916.5 MHz, on-off keyed, or frequency modulated.

30 Long range communications circuitry include a cellular telephone connected to the microprocessor by a modem, or circuitry for employing all digital PCS (personal communications systems), AMPS (advanced mobile phone system), wireless data systems, GEOSAT systems
35 (geosynchronous satellite), or LEOSAT/MEOSAT (low earth

- 42 -

orbit satellite/medium earth orbit satellite). Other voice and data communications systems include Qualcomm's and InMarSat's geosynchronous based service, satellites for low cost messaging similar to two-way paging, CDPD and dedicated, trunked or other UHF and VHF communications.

The external power source can include solar power or fossil fuel power, for example, a vehicle's battery. If fossil fuel power is used, the tag can perform unlimited long distance communications. In embodiments where an external power source is readily available, the battery may be omitted. The tag can also be provided with an RS 232 port for factory testing.

If smaller tags are needed, for example, in the clandestine tracking application 10, of Fig. 9, the battery is omitted or decreased in size. The size of a battery is proportional to the amount of energy it stores. A tag with a battery lasting two years is approximately 30 mm x 30 mm x 5 mm in size. For the clandestine application 10, (shown in Fig. 9), a shorter lived, smaller battery is employed.

Some operations of the microprocessor (such as determining locations, described below) are alternatively performed by a machine performing logical and arithmetic calculations, for example, an ASIC state machine.

Sensors

Sensors 18 (Fig. 2) can include a variety of analog and digital sensors. Each sensor is placed in the communications range of an asset tag 16 or helper tag 20, or is connected by a serial interface to a communications controller in monitoring device 15, described below.

Alternatively, as shown in Fig. 14, the sensor 18 is incorporated into the asset tag 16. The sensor is connected to the microprocessor 100 by a sensor interface 190. If the microprocessor has the capability of

- 43 -

communicating with the sensor directly, the sensor interface is omitted. Similarly, the sensor can be incorporated in a helper tag 20 (Fig. 2), as shown in Fig. 15.

5 Tags communicating with the sensor generally report only deviations in the sensor's input from preset limits. The limits can depend on the absolute time, location of the sensor, or condition of an associated asset (for example, whether a tag has become detached
10 from the asset). Examples of these functions are given below.

 One type of sensor includes a global positioning system (GPS) receiver, for example, a Trimble Navigator, available from *Trimble Navigator, Ltd., 438 East Main*
15 *Rd., IronGate #1, Middleton, RI 02840*. The GPS provides global coverage, continuous operation, and through a GPS receiver supplies position information with an accuracy on the order of 10 meters (or less depending whether the satellites function in a selective availability mode).
20 The receiver also supplies velocity information and a current time. The receiver is typically operated for a few minutes to obtain ephemerides (locations) of the satellites in order to compute a time and position for the tag. This is accomplished using a method similar to
25 that described in Section XA below. Monitoring device 15 (Fig. 1) shuts down the receiver between measurements to save energy.

 Other sensors include temperature sensors, fuel level detectors, brake release detectors and door
30 sensors. These sensors are available from Ivonics, located in Boulder, Colorado, and are incorporated, for example, in the railroad system as described below. These sensors generally communicate with device 15 through an asset tag. Alternatively, the sensors are

- 44 -

incorporated in specialized helper tags that communicate with asset tags.

Other sensors collect audio or video signals relating to an asset, analyze the signals for deviations from a given profile, and transmit compressed versions of the signal when a deviation exists.V. Helper tags

Referring to Fig. 10, a helper tag 20 is illustrated as a rectangular package with dimensions larger than the asset tag 16. Three pins 196, 197, 198 on a face 199 of the helper tag are configured to be inserted in the three recesses 96, 97, 98 on face 99 of the asset tag. In some embodiments, the helper tag includes two sets of quick connect ports 202, 204, 202', 204'. A. Seal Helper Tag

Referring to Fig. 15, one example of a helper tag is a seal 20'. The seal includes a character detector 205 (for example, a shift register) connected between an IR receiver in IR receiver/transmitter 206 and a switch 208. The receiver and transmitter 206 are connected to an asset tag port 103' on face 199 of the seal (Fig. 10).

A second infrared receiver 134' and an infrared transmitter 136' are each placed adjacent to a quick connect port 202', 204' respectively. The second infrared transmitter 136' is coupled to a switch 208 by a timer 210. The timer is connected, through switch 208 to a cellular telephone and antenna 212 by a cellular modem 214 and to a GPS receiver 18 by a sensor interface 190. The seal additionally includes the power controller and battery described above in connection with Figs. 11 to 13. The battery includes a ground connection and a power connection coupled to pins 196, 197, respectively, on face 199 of the seal (Fig. 10). Pin 198 is grounded in this embodiment.

In use, an operator inserts pins 196, 197, 198 on the helper tag in the recesses 96, 97, 98 in the asset

- 45 -

tag. This enables the battery in the seal to supply power to the external power port in the asset tag through pin 197. Assembling the seal and asset tag in this manner also places helper tag port 103 on the asset tag adjacent the asset tag port 103' on the seal, which
5 enables IR communication between the IR transmitters and receivers 134, 136, 206 in the asset tag and seal.

The user next attaches the seal to an asset by threading an aircraft cable through a latch on the asset,
10 and inserting opposite ends of the cable in the quick connect ports 202, 204. The user also threads an optical fiber 216 through the latch and places opposite ends of the fiber in quick connect ports 202', 204'.

The timer causes the second IR transmitter to
15 periodically transmit a pulse along the optical fiber. If the fiber is damaged or broken, the IR receiver will not receive the pulse, and will notify the attached asset tag. The asset tag then notifies another tag, an operations center or a monitoring device over the
20 communications circuitry that the seal may have been damaged.

The character detector in the seal analyzes messages from the attached asset tag, and in response instructs the switch to connect the asset tag port 103'
25 to the appropriate component, for example, the sensor 18 or the modem 214, in the helper tag.

In some embodiments, the seal includes a sensor or actuator that permits the seal to communicate only with a list of selected tags, devices or users. The actuator
30 may include a fingerprint detector or a magnetic stripe reader. Whenever a tag or device not on the list attempts to access the seal, the seal notifies a tag, device or the operations center.

Typically, for added security, a second asset tag
35 and seal combination is placed on the asset, so that if

- 46 -

the first seal is disabled, the second seal may notify other tags or devices. In some embodiments, only one seal includes the GPS receiver and sensor interface.

B. Other embodiments of helper tags

5 In some embodiments, the IR transmitters and receivers in the helper tag and asset tag communicate with a modulated IR, RS 232 or a similar protocol. In addition, the helper tag may include its own microprocessor.

10 The helper tag and asset tag may also be the same size, and be attached through a variety of mechanisms. Alternatively, the asset tag may be larger than the helper tag.

In some embodiments, the character detector in the
15 helper tag checks a special character in the message, indicating whether the helper tag should be placed in a pass through mode or in an operations mode. If the character indicates a pass through mode, the detector causes switch 208 to connect the asset tag port 103' to
20 one communication circuit in the helper tag, for example, cellular modem 214. The helper tag thus retransmits the message received from the asset tag over the communication circuit. In this mode, the seal acts as a communications interface for the asset tag.

25 If the message places the helper tag in an operations mode, the character detector in the helper tag causes switch 208 to couple the asset tag port 103' to all other circuits in the helper tag, for example, the sensor 18, modem 214, IR transmitter 136' and receiver
30 134'. In this mode, the asset tag causes the helper tag to perform specialized functions, including switching the RF communication to a particular gold code. When the helper tag receives the command from the asset tag, the helper tag returns an acknowledgement of the command over
35 the asset tag port 103'. Another function includes

- 47 -

adding a preamble to all signals transmitted over the RF transmitter in the helper tag. The preamble indicates that the signal should be acted on only by a particular asset tag. Similarly, the asset tag may remove a
5 preamble from signals arriving at the RF receiver in the helper tag that specify a particular tag identification number (ID).

The asset tag also causes the helper tag to change the time T over which the synchronizing code is broadcast
10 to a new value. The helper tag subsequently transmits the code for a length of time equal to the new value.

Another embodiment of a helper tag is described below, in connection with the interrogate and respond module. A description of two other helper tags (called a
15 boot and holster) is provided in connection with the monitoring device.

VI. Interrogate and Respond Module

A. Structure

The interrogate and respond module 22 (Fig. 2)
20 includes a helper tag 20, attached to an asset tag 16 in the manner shown in Fig. 10.

Referring to Fig. 16, one embodiment of helper tag 20, (helper tag 20') in the module has a character detector 205 connected to an asset tag port 103' by asset
25 communications circuitry 206'. A switch couples circuitry 206' and detector 205 to local communications circuitry 230 and a radar and voice communications circuitry 232. The local communications circuitry 230 can include an RS 232 link.

30 A power source 110 is connected to the circuitry 206; a microprocessor, by a power control 109. The power source can include a battery or a connection for an external power source. The power control is connected to pins 196, 197 (Fig. 10) and to the local communications
35 circuitry and the radar and voice communications

- 48 -

circuitry. The power control includes the device power control, processor power control and communications power control connected as shown in Figs. 12 and 13. The helper tag is optionally supplied with the detach
5 detector and battery sensors shown in Fig. 13, and a sensor interface as shown in Fig. 14.

Referring to Fig. 17, the voice and radar communications circuitry 232 includes a voice transmitter 240 and a radar receiver 242. The transmitter is
10 connected to a voice synthesizer 244 receiving data from a recirculating FIFO 246. An input of the FIFO is connected to an output of the FIFO by a switch 247. The output of the FIFO is in turn coupled to an input of transmitter 240 by a switch 248. The radar receiver is
15 connected to the synthesizer, the FIFO and the radar transmitter, and provides an interrupt signal to pin 198 on face 99 of the helper tag (Fig. 10).

The asset communications circuitry 206' and the local communications circuitry 230 each include receiver
20 134' and transmitter 136' performing either unmodulated infrared or modulated RS 232 serial communications. B. Operation

In use, the pins 196, 197, 198 in the helper tag are inserted in recesses 96, 97, 98 of an asset tag
25 (Fig. 10). The helper tag supplies power to the asset tag over pins 196, 197, and provides an interrupt signal over pin 198. The helper tag also communicates with the asset tag through the asset tag port 103' placed adjacent the helper tag port 103 on the asset tag.

30 After receiving a message from the asset tag at the circuitry 206' (Fig. 17), the character detector searches for and finds the special character in the message. If the special character indicates pass through mode, the detector connects the switch to the transmitter
35 136' in circuitry 230, and passes the message to that IR

- 49 -

transmitter 136'. The transmitter then sends the message to an appropriate destination. If the local communications circuitry is an RS 232 serial link, the microprocessor in the attached asset tag converts the
5 signal to an RS 232 protocol before transmission. If the character detector determines that the character indicates an operation mode, the detector connects switch 250 to the FIFO. The microprocessor of the asset tag then loads information stored in its memory into the
10 FIFO. (Switch 247 is placed to connect switch 250 to FIFO 246.) The information is derived from communications with the asset tag [or through its local communications]. This typically includes a description of the assets to which other asset tags in the vicinity of the helper tag
15 are attached.

If the radar receiver later receives a standard police radar signal, the radar receiver causes 109 the power control to supply power to the voice synthesizer 244 and transmitter 240, and sends signals activating the
20 synthesizer, FIFO and transmitter. The receiver 242 also supplies an interrupt signal which forwards the signal over pin 198 to the interrupt port in the asset tag. The signal causes the asset tag to halt its communications with any tag other than the helper tag.

25 When activated, the FIFO connects its output to its input through switch 247, and downloads its stored information into the voice synthesizer while simultaneously returning the information to its input port. This allows the FIFO to continue to store the
30 information, even after it has been downloaded into the synthesizer. The information thus remains available at the FIFO the next time the radar receiver is interrogated.

The synthesizer converts the information from the
35 FIFO to voice data and transmits the data to the

- 50 -

transmitter through switch 248. The transmitter broadcasts the data on a police radio frequency (a carrier frequency). The transmitter then connects, through switch 248 to FIFO 246 and transmits the ASCII data stored in the FIFO on the same radio frequency for reception by a suitable computer. Alternatively, the transmitter broadcasts the data on a different radio frequency.

Since the helper tag will receive radar signals only during an emergency (which are expected to be infrequent), transmissions over transmitter 240 may be high power even though the energy source is a modest size.

The information is loaded onto the FIFO generally before the helper tag receives a radar signal. As a result, the helper tag is not required to seek the data during an emergency, when the communication links in the system may be inoperative.

In some embodiments, the radar receiver responds to a radar signal only under certain conditions. For example, if the module is mounted on a vehicle, the receiver will not respond to radar signals received while the vehicle is moving. In this embodiment, it is assumed that the vehicle moves only under normal, non-emergency conditions.

VII. Monitoring Device 15

A. Structure

Referring to Fig. 18, monitoring device 15 includes a personal digital assistant (PDA) 270 coupled to a communications controller 272 by an interface 274. The interface includes a helper tag (a boot) 20' connecting an asset tag 16 to the PDA. A second helper tag (a holster 20" in the interface provides a link between the communications controller and the asset tag. The PDA, for example, is a Newton, available from Apple,

- 51 -

in Cupertino, CA. The Newton has a processor connected to a user interface, a modulated IR communications port and a memory. The user interface allows an operator to access information that has been stored in the memory by
5 the system. The operator also uses the interface to cause the processor to execute stored functions.

The PDA has at least two software applications. The first includes a driver interface for programming tags. In specific applications, the software may include
10 additional functions. For example, in a system where the assets are being transported (for example, application 10, described in Fig. 6 above), the PDA has software for identifying which assets are to be dropped off and picked up at a particular stop, and for reviewing maps. The PDA
15 provides lists for checking off completed stops, and the driver may modify the lists. The PDA also records signatures of the recipients of the assets and verifies the correct assets have been picked up. The second software application monitors incoming messages and
20 routes them to other tags and devices. The messages include information regarding the contents of an asset and messages identifying missing tags and tags with low batteries. If the PDA stores a list of the contents of the assets, it compares the information to the list and
25 notifies other tags or the operator of any discrepancy. The device also assists in downloading the information to an interrogate and respond module. Finally the application includes the location and network formation functional units described below.

30 Boot 20' is a helper tag of the type shown in Fig. 10. Referring to Fig. 19, the boot includes a switch 240 connected to a programming port 297 and a communication port 206', each including an IR transmitter and receiver. The character detector 205 to be described
35 in connection with the helper tag of Fig. 20 is replaced

- 52 -

by a microswitch 298 which senses the presence of an asset tag in the programming port and causes the switch 240 to route signals to the programming port 206'. A manual interrupt 299 provides an interrupt signal to the programming port, which forwards the signal to pin 198 on the boot (Fig. 10). The manual interrupt is activated by pressing a button on a face of the boot. The boot additionally contains a power control (not shown) connecting a battery to pins 196, 197, in the manner described above.

The boot is connected to asset tag 16 (Fig. 18) by placing communication circuitry 206' opposite port 103 of the asset tag (Figs 10 and 11).

The holster 20" is a helper tag of the type shown in Fig. 10. Referring to Fig. 20, the holster 20" has a switch 250 coupled to an asset helper tag port 103 by an unmodulated infrared transmitter 136'. A character detector 205 is interposed between the IR receiver 134' and the switch 250. The switch is also connected to a receiver 300 and transmitter 302 for RS 232 serial communications and to a sensor interface 190 (as shown in the helper tag of Fig. 15). A charge port 304 on the holster is configured to connect the PDA and its attached boot to the holster's power source. Similarly, a power port 306 for the asset tag is connected to pins 196, 197 on the helper tag (Fig. 10). Both ports 304, 306 are connected to an external power port 140' and to the switch 250 by power control 109.

The holster also includes all the features of the asset tag shown in Figs. 11 and 12, but without the RF transmitter and receiver and microprocessor 100. The boot and holster are continually activated by their power controllers. Alternatively, the boot and holster are activated when they receive a message at their respective IR receiver.

- 53 -

When the monitoring device is assembled, the pins 196, 197, 198 in the holster are inserted in the recesses 96, 97, 98 in the asset tag 16, in the manner shown in Fig. 10. The holster supplies power to the asset tag
5 over pins 196, 197, and communicates with the tag through asset tag port 103'. The helper tag (holster 20") is coupled to the communications controller by connecting the RS 232 port (receiver 300 and transmitter 302) to the communications controller.

10 Referring to Figs. 21 and 22, the communications controller includes a PC/104 processor 320 running, for example, Microsoft DOS. The PC/104 (already installed with DOS) is a unit number 6 available from *Tri-M Systems in Coquitlam, British Columbia*. The processor is
15 connected to a random access memory 322, a display controller 324, and a set of RS 232 ports 338. Each port includes an RS 232 transmitter and receiver. The controller is powered by an external power source.

For testing, referring to Fig. 22, the controller
20 324 is equipped with a user interface 323 (for example, a keyboard and monitor) which is removed when the controller is in use. The interface is connected to the display controller 324 in the processor, and managed by display manager software 326 in the processor.

25 The processor is linked to a cellular phone 212' by a cellular modem 214' over an RS 232 link 338. The cellular modem is a ZyXEL U-1496P, available from *ZYXEL/Zero One, 4920 E. La Palma Ave., Anaheim, CA 92807*. The processor includes a communication manager
30 332 that controls the modem. A GPS receiver 339 is coupled to the processor over another RS 232 port 338 and managed by a manager 352. A translator (such as a holster 20") coupled to an RS 232 port 338 converts RS 232 communications from the processor to IR
35 transmissions. A PDA link manager 354 controls the

- 54 -

holster over port 338, causing the holster to communicate with the PDA. A sensor manager 356 controls a port 338 connected to a sensor 18. Other functions in the processor are handled by a master manager 358.

5 B. Operation of monitoring device

(a) Network communications

 An operator employs the monitoring device for programming the asset tags, and for communicating with the asset tags, other devices, sensors and the operations
10 center (Fig. 1).

 Referring to Fig. 18, for example, an operator communicates with a sensor in the system by first selecting an appropriate command on the PDA 270. The PDA then generates a message based on the command,
15 for example, a command for reading sensor data results in the message "PDA #Y to sensor ID # X, send data\." The PDA forwards the message to the boot which forwards the message to its attached asset tag 16 (Fig. 18).

 The attached asset tag checks the destination
20 of the message and determines, for example, that the message must be retransmitted. The asset tag 16 attaches a preamble to the message, indicating "Holster ID # Z, message follows." The holster 20" receives the message from the asset tag, the holster's character detector
25 places the holster in a pass through mode, and forwards the message over the RS 232 link to the communications controller 272.

 The controller determines that it is not the destination of the message and adds the preamble
30 translator ID # W, message follows" to the message. The controller then passes the message to an asset tag attached to an appropriate translator over an RS 232 port.

 The translator asset tag receives the
35 message, strips the preamble, and determines the message

- 55 -

is not destined for it. The translator asset tag passes the message to its IR port and forwards it to a fourth asset tag that is attached to the helper tag containing the sensor. The fourth asset tag determines that the destination of the message is its attached helper tag, and generates a special character that it attaches to the beginning of the message. The helper tag character detector receives the message, detects the character, and connects the switch 208 in the helper tag (Fig. 15) to the sensor interface 190. This allows the fourth asset tag to read data provided by the sensor.

After obtaining the data, the fourth asset tag then formulates the message "from sensor ID #X to PDA #Y, the location is (loc.)." This message is transmitted back to the PDA in a manner analogous to that described above.

In the operations mode, the boot 20' and holster 20" respond to all the commands listed in Section V above, except the boot does not respond to the message commanding a change in the time T over which the synchronization code is transmitted.

Referring to Fig. 22, the processor 320 in the communications controller continuously scans each RS 232 port 338 for incoming messages. The processor routes a message arriving from a particular source to an appropriate manager that examines the destination identification in the message and forwards the message over the appropriate RS 232 port 338. For example, when the controller receives a message addressed to the PDA over its cellular telephone from the operations center, the manager 332 forwards the message to the PDA link manager 354, and in turn to the holster 20" for transmission to the PDA.

The processor buffers the incoming messages in memory when it becomes necessary to match the

- 56 -

communication rate of a transmitter with the communication rate of a receiver in a particular RS 232 port. Certain messages are buffered until the PDA queries the controller for the message. For example, the controller may store messages arriving from the operations center until interrogated by the PDA.

The boot and holster also receive RF messages directly from other electronic tags or sensors in the system. Similarly, the asset tag 16 communicates directly with other asset tags.

(b) Programming asset tags

Referring to Fig. 18, the operator employs the boot 20' to store information directly on an asset tag 16," for example, software for maintenance and upgrades, or to store information relating to an asset. The operator inserts pins 196, 197, 198 on the boot into recesses 96, 97, 98 on the tag 16." This places the programming port of the helper tag adjacent to port 103 on the asset tag, to enable communication between the ports.

Inserting the pins of the boot into the recesses of the asset tag causes the microswitch 298 (Fig. 19) to activate the switch 240, and connect the programming port 297 to the microprocessor in the asset tag. The switching action simultaneously disconnects the communication port 206' from the microprocessor.

The operator presses the button on the boot to activate the manual interrupt 299. The boot transmits the interrupt over pin 198 to the interrupt port on the asset tag 16," causing the attached tag to discontinue all communications with other devices and look for information from its helper tag port 103.

The operator then selects information on the PDA through the user interface. The PDA transmits the information to the boot 20', which retransmits the

- 57 -

information over the programming port to the IR receiver in the helper tag port to the attached asset tag. The microprocessor 100 in the tag receives the information and stores it in memory.

5 The operator then removes the asset tag from the programming port. This action causes microswitch 298 to cause switch 240 to disconnect from the programming port and connect to the communication port. If
10 the PDA is battery powered, it is recharged by inserting a power port on the PDA in the PDA charge port 304 of the holster (Fig. 20).

VIII. Operations Center

The operations center 13 (Fig. 1) includes one or more workstations linked together by a network. For
15 example, the center includes Pentium PCs and Sun Sparkstations linked by a 10Base T Ethernet network. The workstations are powered by an uninterruptable power supply. Commercial power is backed up by a generator.

At least one of the workstations has a cellular
20 communications modem to land lines. In addition, redundant equipment is provided to enable operation when a workstation or the communications interface fails. The workstations include database software that is tailored to the system, for example, Access or Oracle software,
25 available from Oracle Corporation, Federal Group, 3 Bethesda Metro Center, Suite 1400, Bethesda, MD 02814. The workstations also include geographic visualization software with raster or vector maps tailored to the area of operation, for example Blue Marble or available from
30 *Blue Marble Geographics, 46 Water St., Gardner, ME 04345.*

Accounting software in the operations center tracks billings and inventory control software manages inventory lists stored in the center. Communication
35 software controls communication with the system, and

- 58 -

client communication software manages communications with clients, as described below. Custom communications gateway software controls communications between the center and vehicles carrying assets. If the assets are stationary, the center is directly connected to particular asset tags, for example, a gateway tag. This is done by using an RS232 connection with standard protocol.

IX. Detailed description of the asset tag microprocessor

10 A. Overview

The processor memory includes entity, gateway and data tables. The entity table includes a list of identification numbers of tags, devices or sensors to which the tag will respond. The numbers are stored in the identity store of each tag. The gateway table includes a list of gateway tags with which the tag communicates. The data table stores information sent to the tag.

Referring to Fig. 23, each tag's microprocessor software 100 includes a hierarchy of functional units. The functional units, from lowest to highest in the hierarchy, include self-assessment 370, joining 372, location 374, network formation 376 and network use 378 units. Certain units are omitted from the tag in specific applications, as described below.

Each unit is in communication with the units adjacent to it in the hierarchy. In addition, a feedback path 380 is provided between the network formation unit and the joining unit.

30 A communications router 382 couples each unit to an ISO 7 protocol stack 384. The stack controls the IR transmitter and receiver and RF transmitter and receiver in the asset tag.

Each functional unit is standardized in its communication with the router and with its neighboring

- 59 -

functions in the hierarchy. This allows each unit to be easily replaced with a new unit.

Each unit is active so long as the unit below it in the hierarchy is active. Once a lower unit becomes
5 inactive, all the units above it in the hierarchy become inactive. Other devices that automatically network with the tags, for example, monitoring devices and sensors, typically also include all the functional units shown in Fig. 23.

10 B. The router

Each unit controls the router by indicating whether a communication should occur over the IR transmitter or the RF transmitter in the tag. If RF is chosen, the unit also selects the frequency, time slot
15 and code to be used in the transmission. In certain applications, a unit may communicate directly with the transmitter or receiver by bypassing the router (for example, the intertag distance subunit, described below).

When the router receives a message from the IR or
20 RF receiver, the router sends the message to the network use unit 378, which checks whether the message is addressed to the tag. This is done by comparing the tag's identification number to the number in the message. If it was not addressed to the tag, the router ignores
25 the message. If the message was properly addressed, the unit determines whether the final destination of the message is the tag itself, an attached helper tag, or another component.

If the message is for the tag itself, the network
30 use unit acts on the message. The commands obeyed by the unit include the commands performed by the helper tag in the operations mode, listed in Section V.

If the message is for an attached helper tag, the unit adds a special character to the beginning of the
35 message, and a different special character to the end of

- 60 -

the message. The special characters are excluded from a set of characters used for transmitting messages. The router then sends the message to the helper tag. The first special character causes the helper tag to switch
5 from a pass through to an operations mode. The second special character causes the helper tag to switch back to a pass through mode.

If the message is for another component, the router sends the message to the component through an
10 appropriate transmitter. If the router employs the RF transmitter, it first transmits the synchronizing code over the period T to allow other asset tags enough time to synchronize to the signal and thus receive the code.

A detailed description of each functional unit in
15 the processor follows.

C. Tag functional units

(a) Self-assessment

Referring to Fig. 24, the self assessment unit 370 includes a battery subfunction 400, an active
20 subfunction 402 and an attachment subfunction 404, arranged in increasing levels in the hierarchy.

Referring to Fig. 26, when the tag is active (step 429), for example, after an interrupt has awakened the tag from its comatose state (Section III), the
25 battery subfunction monitors the state of the battery through the battery sensor 166 (Fig. 13) (step 430). If the battery is low (step 432), the subfunction notifies the active subfunction (step 434). The active subfunction in turn passes this information to higher
30 levels of the hierarchy.

The battery subfunction also notifies a gateway tag through the router that the battery is low (step 434). The message includes the tag's identification number. The subfunction then sets the
35 tag's timer to a maximum value (step 436) and decides

- 61 -

whether to deactivate the tag (step 438). If the subfunction does not deactivate the tag, it returns to step 420 and repeats the process.

Referring to Fig. 27, the active subfunction
5 monitors the interrupt port 107 (Fig. 11) of the tag for an interrupt signal (step 450). The signal is typically supplied by a boot in the monitoring device when the tag is attached to the programming port of the boot. When the active subfunction detects the interrupt, it informs
10 the processor power control to increase the power supplied to the microprocessor (step 452) and the IR transmitter and receiver (step 454). This causes the tag to leave the comatose state. The subfunction then resets a pointer to a first unit of available memory in the
15 entity and gateway tables (step 456).

The active subfunction also responds to the following commands received from the router. The subfunction replies to a query regarding the tag's ID. Another command indicates to the subfunction that it will
20 be commanded by a particular entity. If the entity ID is not recorded in the entity table, the subfunction records the ID in the table. Another command causes the subfunction to clear certain data in the data table and write new data. This can include data relating to the
25 asset's shipper, destination etc. Other messages cause the subfunction to concatenate new data to the data table. The subfunction checks for and notifies another asset tag or the monitoring device 15 (Fig. 1) of any data overflow from the router.

30 Still another command causes the subfunction to read and communicate the data in the data table, the data in the gateway table, and the information in the entity table. The subfunction also deletes any erroneous or unwanted information, such as an entity ID, and

- 62 -

activates and uses the RF transmitter and receiver in response to appropriate commands.

The active subfunction informs the attachment subfunction of the tag's status. For example, the active subfunction can send the message that the tag is inactive, as commanded by a particular entity. The attachment unit verifies that the entity is listed in the entity table. If it is, the attachment subfunction becomes inactive.

10 Referring to Fig. 28, when the tag is active (step 460), the attachment unit verifies the detach detector 164 (Fig. 13). When it determines that the tag is detached (step 462), it waits 60 seconds to enable the operator to attach it to an asset (step 463). If still
15 detached, the subfunction informs each gateway in the gateway table through the router (step 464). The subfunction continues to transmit this message until it receives notification that the tag ID and its detachment have been recorded by one or more entities (step 466).
20 If one of the entities is listed in the entity table (step 468), the attachment unit becomes inactive (step 469). The tag also notifies the joining unit of the detachment in step 464. Alternatively, the subfunction may cause the router to notify a device, such as a user-
25 operated device, instead of a gateway of the detachment.

(b) Joining

The joining functional unit 372 includes an existence subfunction 470 connected below a hierarchy subfunction 472 (Fig. 24).

30 Referring to Fig. 29, when the attachment subfunction indicates to the existence subfunction that the tag is active (step 480), the existence subfunction causes the tag to look for other tags by attempting to initiate communications through the router on a
35 particular hailing channel (482). If other tags exist

- 63 -

(step 484), the subfunction notifies the hierarchy and supplies the identity codes of the tags (step 486).

If a group of tags are located near the asset tag, the subfunction chooses one asset tag (a netmaster) to perform some of the location and network formation units. This prevents the tags from duplicating calculations already performed by other tags.

Referring to Fig. 30, the hierarchy subfunction communicates with the other tags to determine if the other tags have designated a particular tag as netmaster (step 487). If the tag locates a netmaster (step 488), it notifies the network formation unit through the location unit (step 489). Otherwise, the tag chooses the tag with the highest identification number as the netmaster, and chooses the tag with the next highest identification number as the assistant netmaster (step 490).

The netmaster periodically surveys the tags for new or missing members, and notifies the hierarchy subfunction. The netmaster uses a particular hailing frequency for its communications. If the netmaster disappears or indicates its battery is low, the assistant netmaster replaces the netmaster, selects another assistant, and informs the hierarchy subfunction of the change.

25 (c) Location

Referring to Fig. 25, the location unit 374 determines the locations of tags in the network using the following subfunctions: intertag distance 500, verify range 502, refine range 504, location compute 506 and detect change 508 (Fig. 25).

Referring to Fig. 31, the intertag distance subfunction 500 determines the distance and the covariance (or measurement error) of the distance between the tag and each tag identified by the joining unit (step 530). The subfunction also communicates with the tags to

- 64 -

obtain location information for the tags (step 532). The location information is typically obtained from a GPS sensor located near the tags. The subfunction organizes this information in the tag memory (step 534), and
5 notifies the verify range subfunction 502 that the information is available (step 536).

The verify range subfunction verifies these ranges, and eliminates ranges that are clearly erroneous. The refine range subfunction provides accurate estimates
10 for the ranges eliminated in the previous subfunction. The compute locations subfunction transforms the corrected ranges supplied by the refine range function into an arbitrary set of coordinates for each tag. Using the location information, the location subfunction
15 transforms the arbitrary coordinate system into an absolute coordinate system, and determines a location for the tag.

The detect change subfunction receives the location information from the verify range subfunction,
20 and monitors the ranges of other tags in the network to detect a change in the tag's location.

Each of these subfunctions is discussed in more detail below.

(d) Network Formation

25 The network formation unit includes a determine topology subfunction 560, a sentinel subfunction 562 and a message subfunction 564 (Fig. 25). The determine topology subfunction receives information from the change detection subfunction indicating changes
30 in the existence of a tag, and the location of the tag over time. Using this information, the subfunction determines a set of neighbors for each tag. A topology description of the set is found using a Delaunay triangulation, described below, which describes the
35 topology as a convex hull and identifies the set of

- 65 -

nearest neighbors. When the determine topology subfunction receives updated information from the location unit, it modifies the topology description.

The sentinel subfunction identifies tags that are located on the convex hull. These tags (referred to as sentinels) are assigned to communicate with tags outside the hull. The subfunction also identifies gateway tags capable of long distance communication and PDAs acting as a man-machine interface. The subfunction stores the identity of each gateway in the gateway table.

The sentinel subfunction communicates the identity of the sentinels to the hierarchy subfunction 472 (Fig. 24) over the feedback path 380. The hierarchy subfunction causes the sentinel to search for new tags, extending the range of the netmaster.

The sentinel subfunction also responds to commands received over the router 382 (Fig. 23). Each time a command from an entity is routed to the sentinel subfunction, the subfunction first verifies that the source of the message is listed in its entity table. If the entity is not listed, the subfunction ignores the message. If it is listed, the unit responds to the command.

The first command received by the sentinel subfunction queries whether a tag with a given identification exists. The subfunction checks whether any stored identification matches any stored tag's identification number. If it does, it answers "yes" to the entity. The unit then checks whether the entity is listed as a gateway in the gateway table. If it is not listed, the unit records the identity of the entity as a gateway. If there is insufficient memory to store the identity in the table, the unit communicates with the oldest recorded gateway in the list. If it is unable to establish communication, it deletes the gateway from the

- 66 -

list. The unit repeats this for each gateway in the list.

The second command asks whether there are tags which have not been queried by a gateway. The
5 subfunction looks at its gateway table, and if the gateway ID is not listed, the unit replies with the tag ID.

The last command asks whether a tag with a particular ID has registered the gateway. If the tag ID
10 is the same as that in the message, the unit checks its entity and gateway table and records the ID of the gateway if it is not already present in the table. If there is insufficient memory to store the gateway, the tag repeats the process described above to delete non-
15 communicating gateways.

The propagate message subfunction defines a network containing a set of efficient communication links between tags using the topology derived by the topology subfunction. The subfunction accomplishes this by
20 assigning tags in the networks as relay points for communications between a pair of tags.

The subfunction also chooses a second netmaster based on the topology. The second netmaster redefines the network whenever a tag is added or removed
25 from the network, and communicates information regarding the changes in the network to the gateways listed in the gateway table. In some embodiments, the netmaster only notifies the gateway of changes regarding specific tags. The other tags in the network do not duplicate the
30 functions of the netmaster.

The netmaster polls members of the network periodically to make sure they still exist. Each tag in the network checks the time interval since it was last polled by the netmaster. If the time interval exceeds a

- 67 -

threshold, the network formation unit in the tag removes the tag from the network.

When the network grows too large, the network formation subfunction divides the network into smaller subnetworks. The subfunction designates channels of communication within a subnetwork that do not interfere with communications in the other subnetwork. In addition, communications within a subnetwork are coded in compacted form. For example, the subfunction assigns each tag a truncated ID. If spread spectrum communications are used, the subfunction uses a short code for communication relayed between tags.

The subfunction designates a tag for maintaining communication with another subnetwork. When the subnetwork grows too small, the subfunction queries the tag for the existence of the other subnetwork. If the subnetwork still exists, the subfunction merges the two subnetworks.

(e) Network Use The network use unit includes a manager subfunction 590, an agent subfunction 592, an inventory subfunction 594, and a read and write subfunction 596 (Fig. 25).

Referring to Fig. 32, the manager subfunction continually checks whether route information is available in the network (step 608). The information is typically provided by a monitoring device, computer or operations center.

When the information is available, the subfunction stores information regarding the transportation route scheduled for an asset attached to the tag and a checking frequency (step 610). The subfunction then surveys the network to determine whether there has been an update to the route information (step 611). If there is, the subfunction returns to step 610 and stores the new information.

- 68 -

If no update information is available, the subfunction examines the location of the tag (as supplied by the location unit) (step 612), and compares the location to the route (step 614). If the location of the asset deviates from the route (step 616), the subfunction informs a gateway through the network (step 618) and returns to step 611. If no deviation exists, the subfunction repeats the process by returning to step 611.

As shown in Fig. 33, the agent subfunction 592 stores information regarding the destination of the attached asset (step 630). When a deviation from the route is reported by the manager subfunction (step 632), the agent subfunction queries the network regarding available alternative routes (step 634). If a more effective alternative route is available (step 636), the subfunction notifies a gateway (step 638).

In some embodiments, the route includes a series of locations where the conveyance carrying the asset is changed (transshipment points). The subfunction seeks a conveyance for the next part of the route at each transshipment point. The subfunction also searches for other tags following a similar route, and searches for a proper conveyance for carrying all the assets together. The subfunction then informs a tag on the desired conveyance that the conveyance should load the assets. The subfunction thus performs functions analogous to a node in a packet switched network.

The inventory subfunction maintains a list of assets in the network. This subfunction is particularly useful in a warehouse application.

The read and write subfunction stores and retrieves data from the asset tag's memory. The subfunction is activated by the following commands from one or more entities stored in the entity list.

- 69 -

1. Read data to one or more identified entities.
2. Clear data as instructed by one or more identified entities, and write certain data.
- 5 3. Append the following data as instructed by one or more identified entities.
4. Add an identified entity to the entity table as instructed by one or more entities.
5. Delete an entity from the entity table.
- 10 6. Read the entity table.
7. Read the gateway table.

X. Detailed design of the location unit

A. Intertag Distances Subfunction

(a) Baseline Method.

15 The intertag distances subfunction (Fig. 34) in a tag 640 determines the tag's distance, that is, range, from another tag 642 by measuring the time required for an electromagnetic pulse to travel round-trip between the two tags. There are several methods for
20 measuring this time. A baseline method is described first, followed by descriptions of alternate methods. Most of the methods require that the pulse-originator tag have a cooperating transponder tag for a reference, with one clock on each tag. However, the last alternate
25 method describes a dual-clock technique wherein both clocks are on the originator tag - the other tag is simply a repeater or a repeater with a delay. In fact, if a delay is not required, another active tag is completely unnecessary in the last alternate method; the
30 reference can be totally passive, for example, a mirror or a three-dimensional corner reflector.

The resolution of the range measurement is proportional to the time-measuring resolution of the tag, so it is advantageous to use methods of range measurement

- 70 -

that effectively increase the time-measuring resolution. The methods described below take advantage of the tag's ability to detect a pulse received during a period of the tag's time-measuring clock. For a single pulse, the tag
5 can detect that the pulse arrived during the period but not when during the period. Therefore, the range-measurement methods use multiple pulses and multiple successive measurements. In the baseline method, small offset delays of differing but known lengths are
10 inserted, one per round-trip, in the round-trips of the pulses, to stagger the arrivals of the pulses. Thus, for a set of ranges, the offset delays cause a fraction of the pulses to return in a later time-measuring clock period in any particular range measurement. As described
15 below, the method determines the range from this fraction.

In the baseline method, the originator tag and the transponder tag initially have identical clock periods. To insert a multiple of an offset delay "d" in
20 the round-trips (Fig. 35), the originator alters its clock period slightly. The length of the alteration is an amount of time which provides a one-clock-period difference with the transponder's clock over a measurement interval. The measurement interval is an
25 amount of time equal to the number of pulses sent, "N," multiplied by the time interval between the sending of each pulse. Although the originator's clock may be altered above or below the frequency of the transponder's clock, the embodiment of the baseline method described
30 herein uses an originator's clock with a higher frequency, that is, a shorter period, than the transponder's clock.

The difference in clock periods results in the insertion of differing but known delays in the round-
35 trip time, in the following way. Like the originator

- 71 -

tag, the transponder tag can detect that a pulse was received during a period of the transponder's clock. The transponder participates in the round-trips by resending the pulse received from the originator. A delay in
5 resending occurs because the transponder does not detect, and thus cannot resend, a pulse until the end of the period of the transponder clock during which the pulse is actually received. If the originator and the transponder use identical clock periods, pulses sent regularly
10 according to the originator's clock always arrive at the same point within a period of the transponder's clock. With identical clock periods, such pulses are also always re-sent at the same time after arrival, that is, with the same "delay," adding nothing to range measurement
15 resolution. A transponder with a slightly longer clock period detects and re-sends pulses with differing delays, because each regularly-sent pulse does not arrive at the transponder at the same time within a period of the transponder's clock.

20 In the baseline method (Fig. 37), the originator tag transmits "N" pulses (step 650) based on its clock such that the time between transmitted pulses, measured in clock pulses and designated "spulses," is constant. In the preferred embodiment, the difference in
25 clock periods is such that after sending "N" pulses, the higher frequency clock has gained an entire clock period compared to the lower frequency clock. After a range delay, which is the time the pulse takes to reach the transponder and which is designated "range" in Fig. 35,
30 each pulse is detected by, or clocked into, the transponder (step 652) as shown in Fig. 35 by the vertical arrows designated "transponder detections." The pulses are detected only at the end of the transponder clock period in which they are received. Although
35 Fig. 35 shows the transponder's detections occurring in

- 72 -

order from earliest after a pulse is actually received to latest, the method does not require and generally will not have this ordering. Depending on the range and the relative timing of the originator and transponder clocks, the cycle of differing delays can begin anywhere. The earliest pulse has an ambiguity "x" which has a maximum actual value of an offset delay "d." Ambiguity "x" refers to the interval between the pulse's arrival at the transponder and its actual detection by the transponder at the beginning of the next transponder clock period. The offset delay "d" is equal to the difference between the transponder clock period and the originator clock period multiplied by the number of originator clocks between the sending of pulses, i.e., "spulses." In accordance with this illustrated embodiment of the invention, the offset delay "d" is also equal to the period of the transponder clock divided by the number of pulses "N" used in the measurement. Each successive transponder clock period is additionally delayed, relative to the time the originator pulse is received, by a multiple of the offset delay "d."

In this baseline method, the transponder can also insert an extra delay, designated "tdelay," (step 654 in Fig. 37) sufficiently long such that the pair of tags can operate in a half-duplex fashion, that is, such that the originator has sufficient time after sending a pulse to switch to attempting to detect the transponder's pulse's return. This extra delay is unnecessary if full-duplex operation is possible, that is, if pulses can be re-transmitted by the transponder and detected by the originator using an independent return path, such as a different electromagnetic frequency. After the extra delay, the transponder transmits a pulse back to the measuring device (step 656). After the range delay, the pulse arrives back at the originator. Since the phase

- 73 -

relationship between the two clocks is arbitrary, there is further created another ambiguity "y" (which also has an actual value which may differ by an amount up to "d") (Fig. 35) between the arrival time and the time that the pulse is actually detected at the originator (step 658).

For each individual measurement, the ambiguity "x" increases by " $(n-1)d$ " over the initial measurement, where "n" is the measurement number in Fig. 35. The ambiguity "y" decreases by the same amount, taking on negative values in some measurements such as 9 and 10, such that " $(x + y)$ " remains a constant having a maximum possible value that is " $2d$ " greater than its minimum possible value. This constant is a result of the fact that " t_d " (the total time to detection of the returning pulse) is the sum of "y," two range delays, " t_{delay} ," and "x." Because " t_d ," the range delay, and " t_{delay} " are constants, "y" must decrease by the same amount that "x" grows, and vice versa.

Whenever a returning pulse arrives at the originator too late to be detected by a clock time " t_d ," that pulse is recognized by the next clock " t_s " time units later (" t_s " is the originator clock period). Due to the relationships between the clocks, described above, there are never more than two such clock periods involved in a particular measurement. From the number of measurements which fall into each of the two clock periods before and after time " t_d ," the range is determined to a resolution which is approximately $1/N$ of the originator's clock period.

In particular, if the range is such that all pulses arrive back at the originator within the same clock period in each measurement (Fig. 36), the range is readily determinable. In this case, all pulses will have had a round-trip time of " t_d ." The average round-trip

- 74 -

range delay is determined by subtracting from "td" the extra delay, "tdelay," and the constant (x + y):

$$\text{round-triprange delay} = td - \text{tdelay} - (x+y) \quad (1)$$

In this case, where all pulses arrive back within the same clock period, the constant (x+y) can take any value
5 from (ts - d) to (ts+d). The constant's actual value within these bounds is unknown because delay "d" corresponds to the resolution of the measurement method. Thus the actual range delay falls within maximum and minimum round-trip range delays given as:

$$\text{round-triprange delay}_{\max} = td - \text{tdelay} - ts + d \quad (2)$$

$$\text{round-triprange delay}_{\min} = td - \text{tdelay} - ts - d \quad (3)$$

10 Thus the range is half the round-trip range delay multiplied by the speed of light "c":

$$\text{range}_{\max} = \frac{c(td - \text{tdelay} - ts + d)}{2} \quad (4)$$

$$\text{range}_{\min} = \frac{c(td - \text{tdelay} - ts - d)}{2} \quad (5)$$

If the range is great enough such that "n" of the "N" measurements slip to the next clock period as shown in Fig. 35, (n)(ts)/N is added to the round-trip range
15 delay:

$$\text{round-triprange delay}_{\max} = td - \text{tdelay} - ts + d + \frac{(n)(ts)}{N} \quad (6)$$

- 75 -

$$\text{round-triprangedelay}_{\min} = td - t\text{delay} - ts - d + \frac{(n)(ts)}{N} \quad (7)$$

and the range is:

$$\text{range}_{\max} = \frac{c(td - t\text{delay} - ts + d + \frac{(n)(ts)}{N})}{2} \quad (8)$$

$$\text{range}_{\min} = \frac{c(td - t\text{delay} - ts - d + \frac{(n)(ts)}{N})}{2} \quad (9)$$

The average of the receiving times of all the pulses is $td + (n/N)(ts)$. Therefore, the range calculation reduces to

$$\text{range}_{\max} = \frac{c(\frac{(\text{Totalclockcount})(ts)}{N} - t\text{delay} - ts + d)}{2} \quad (10)$$

$$\text{range}_{\min} = \frac{c(\frac{(\text{Totalclockcount})(ts)}{N} - t\text{delay} - ts - d)}{2} \quad (11)$$

5 where Totalclockcount is the sum of all the clock counts used in the measurement and $(\text{Totalclockcount})(ts)$ is the sum of all the receiving times, as illustrated below. For example, as shown in Fig. 35, if seven measurements are detected at time "td" and three are detected at time
 10 "td+ts," the sum of the all the receiving times is $7(td) + 3(td+ts)$. Dividing the sum of all the receiving times by "N" provides a time measurement of resolution ts/N , effectively.

Turning now to a discussion of reducing
 15 errors in the method, the range calculation method is refined further as necessary. The difference between the minimum range value and the maximum range value is $(c)(d)$. In the method, a reference range value is taken

- 76 -

to be the average of the maximum range value and the minimum range value. Thus, the actual range value is always within $(c)(d)/2$ of the reference range value, that is, $(c)(d)/2$ is the "primary range error" with respect to
5 the reference range value.

The range value is considered "nominal," meaning within a minimum of error, when the group of transponder detections is centered in a "ts" window 3500 as illustrated in Fig. 35. In other words, when the
10 actual range produces centered transponder detections, the reference range value accurately represents the range. However, if the actual range is slightly shorter, the group of detections appears slightly shifted towards the left end 3502 of the window, but the range
15 measurement remains the same. Similarly, if the actual range is slightly longer, the group of detections appears slightly shifted towards the right end 3504 of the window, but the measurement remains the same. In addition, the same observations apply to shifts
20 corresponding to the detections of the returns of the pulses. That the measurement remains the same for different ranges amounts to measurement error. This measurement error corresponds to the primary range error noted above.

25 With respect to Fig. 38, a simple example is now used to illustrate the primary range error as a function of range. In the example, an originator tag is stationary and a transponder tag is mobile. At first the transponder tag is at an actual range 3702 from the
30 originator tag such that the primary range error is minimal, that is, the measured range value is nominal and the transponder detections are centered. The transponder is then moved slightly further from the originator tag, corresponding to a slightly larger range 3704. This
35 larger range 3704 generates the same range measurement

- 77 -

value as the original nominal range 3702, but this measurement value is now low compared to the actual range value, resulting in a negative error level 3706. As the range is increased further, the magnitude of the negative error level also increases. At a range 3710 midway in-between the original actual range 3702 and an actual range 3708 corresponding to the measurement method's next highest nominal range value, the error switches from negative to positive. The switch occurs because at this point the measurement method reports as a value, the range 3708, which is greater than the actual range 3710. If the transponder's actual range is then increased further, the magnitude of the error decreases as the transponder tag's actual range closes in on the new nominal range 3708. Thus, as a function of range, the primary error forms the sawtooth pattern shown in Fig. 38.

With respect to phase-related errors (Fig. 35), the relationship described above between the transponder clock period "tt" and the initiator clock period "ts" can be described in terms of the "N" pulses sent at "spulses" intervals, as follows.

$$ts = tt \frac{((N)(spulses) - 1)}{((N)(spulses))} \quad (12)$$

The incremental offset delay "d" between the transponder's periods, in terms of the initiator period "ts," is $(tt - ts)(spulses)$. Because there are $(N - 1)$ such delays "d" plus a leading delay, "x," which cannot be larger than "d," the largest possible interval between the pulse's arrival at the transponder and the transponder's detection of the pulse is $(N)(tt - ts)(spulses)$, or, substituting for "ts" from the previous equation, the largest arrival-detection interval is "tt." Thus, a detection occurring "tt" after arrival is referred to as a "maximally-late" detection.

- 78 -

The "maximally-late" detection can cause a bias in the range measurement. For example, as shown in Fig. 36, it is possible that a phase-related error could cause the tenth transmission, which should be detected at an earlier time 3504, to be detected instead at a later time 3506 (not shown to scale). In this example, this "maximally-late" detection does cause a bias in the range measurement.

The bias occurs because the maximally-late detection provides a maximally-late re-transmission which, due to "tt" being larger than "ts" in this illustrated embodiment, causes the return of the pulse at the originator 3508 to be detected one originator clock period late. In other words, in this particular "maximally-late" situation, the return of the pulse 3508 is extended past the originator clock 3510 and therefore is not detected until the next originator clock 3512 even though it should have been detected at clock 3510.

Given that the length of the "ts" window is $tt - tt/(N) \text{ (spulses)}$ (from the previous equation), the maximally-late interval 3514 is $tt/((N) \text{ (spulses)})$ time units in length. With the clock phasing distributed uniformly, the probability of the re-transmission occurring at the maximally-late time is the interval 3514 divided by the total time in which the transponder clock period can exist, or "tt":

$$\frac{\left(\frac{tt}{((N) \text{ (spulses)})} \right)}{tt} \Rightarrow \frac{1}{((N) \text{ (spulses)})} \quad (13)$$

As described below, (N) (spulses) is typically in the tens of thousands for these tags, so the probability of the maximally-late retransmission is less than 0.01%. In addition, when the error occurs, it adds a maximum range delay error of ts/N , because the detection one originator

- 79 -

clock period "ts" late is averaged across all "N" measurements in the baseline range measurement method.

Furthermore, the latest detection-re-transmission is the only one that can cause such an error, because, as discussed below, the second-latest transponder detection/re-transmission is always in the "ts" window. The arrival-detection interval corresponding to the second latest detection/re-transmission is one "d" delay shorter than the interval for the latest, or $(N-1)(tt-ts)(spulses)$. Subtracting the second latest's interval from the length of the "ts" window, "ts," and substituting for "ts" yields a difference of $(tt/N)(1-1/spulses)$. As a result of "spulses" being large and positive, the second term in $(1-1/spulses)$ is small and less than one. The total result is positive, regardless of the chosen values of "N" or "spulses," which means that the second-latest re-transmission is always in the "ts" window. Thus, for most practical purposes, the baseline method is insensitive to phasing between the two clocks. However, this secondary error due to phasing does have the effect of increasing the apparent primary range error, because the resulting additional round-trip range delay, which is equal to "ts," makes the range seem longer. With reference to Fig. 39, this secondary error, when it occurs, creates an altered "jump" ΔR in the apparent range. The jump ΔR occurs closer in range by a distance ΔR of $(c)(ts)/2$, which is the speed of light multiplied by half the additional range delay. This effectively offsets the primary error's sawtooth pattern by an error increase ΔR of $(c)(ts)/(2)$. This error increase ΔR is corrected, on average, by subtracting the error increase ΔR , discounted by its probability, from the range calculation:

35

(14)

$$range_{max} = \frac{c \left(\frac{(Totalclockcount)(ts)}{N} - tdelay - ts + d \right)}{2} - \frac{(c)(ts)}{(2(N)(spulse)}$$

$$range_{min} = \frac{c \left(\frac{(Totalclockcount)(ts)}{N} - tdelay - ts - d \right)}{2} - \frac{(c)(ts)}{(2(N)(spulses)}$$

(1)

5)

If the expression (tt-ts)(spulses) is substituted for d, an average range_{aver} of range_{max} and range_{min} can be calculated as follows. For a fixed calculation, "N," "ts," "tt," and "spulses" are constants. In the baseline method, "tdelay" is also a constant, which simplifies the calculation to:

$$range_{aver} = \frac{c \left(\frac{(Totalclockcount)(ts)}{N} - Constant \right)}{2} \quad (16)$$

$$Constant = tdelay - ts - \frac{ts}{(N)(spulses)} \quad (17)$$

To simplify computation, the calculation of the range is reduced even further to:

- 81 -

$$\text{range}_{\text{aver}} = \text{Constant1}(\text{Totalclockcount}) - \text{Constant2} \quad (18)$$

$$\text{Error1} = \frac{c(tt-ts)(\text{spulses})}{2} \quad (19)$$

$$\text{Constant1} = \frac{(c)(ts)}{2N} \quad (20)$$

$$\text{Constant2} = \frac{c(\text{Constant})}{2} \quad (21)$$

where $\text{range}_{\text{aver}}$ is within Error1 of $\text{range}_{\text{max}}$ and $\text{range}_{\text{min}}$.

The same principles apply if the clock frequencies are reversed, that is, the originator clock has a lower frequency than the transponder clock. In this case, all transponder detections are contained within the clock period of the measuring device. However, because the "ts" window is never quite filled with transponder detections, the group of transponder detections has a range error of the same probability and magnitude described above.

The overall resolution of the range calculation is Constant1 or $(c)(ts)/(2N)$, which is the slope of the range calculation equation. In other words, every increment in Totalclockcount corresponds to another $(c)(ts)/(2N)$ of range. Thus, for example, for "N" on the order of 100, the resolution is on the order of 1% of $(c)(ts)$.

This resolution enhancement allows the baseline method to measure ranges using a relatively low bandwidth. Transponder and initiator clock frequencies on the order of 20 MHz (a 50 nsec clock

- 82 -

period), combined with "N" = 100, produce a range delay measurement resolution on the order of 0.5 nsec with \pm 0.25 nsec maximum error. Higher values of "N" and higher clock frequencies result in finer resolution.

5 At ten microseconds between pulses, with a 20 MHz clock, "spulses" = 200. If "N" = 100, the frequency ratio of transponder clock to originator clock is 1.00005, that is, 50 ppm (parts per million). At 20 MHz, this yields a 1,000 Hz difference between the transponder
10 clock frequency and the originator clock frequency.

 Of course, secondary errors external to the baseline method itself can affect the resolution of the method. The transponder and originator clocks each have frequency error rates on the order of 1-10 ppm when
15 operating at the same temperature. The frequency errors are somewhat larger if the temperatures are substantially different. Thus, a working clock rate of 20 MHz has a maximum error of 200 Hz for each clock. If the relationship between the originator clock and the
20 transponder clock requires a frequency difference of 1000 Hz, and both clocks are in error by the maximum amount and in opposite directions, the difference error is 400 Hz, which is significant in relation to the frequency difference. In addition, range measurement errors due to
25 clock frequency error increase as the range increases. However, if the delay times inserted by the measuring device and the transponder are large compared to the range being measured, these errors are nearly constant. If the delay times are not large, or if clocks with lower
30 error rates cannot be used or higher frequencies are needed, an alternate clock method described below is used.

 Transmission of pulses can also result in group delays and pulse smoothing such that, for example,
35 the pulses are delayed and have slow rise times. To

- 83 -

compensate, the group delays are measured and added to "tdelay" in the range calculation.

In addition, where the pulses are received by a threshold detector, a fluctuation in the perceived
5 amplitude of the pulses at the detector can result in detecting the pulse with an erroneous clock count. With "N" = 100, such fluctuations or variations in the perceived amplitude are controlled to less than 1% to achieve reliable measurements. If "N" is increased, the
10 fluctuation control is tightened or the transmission bandwidth increased or both. Any variations in the clock periods are also controlled to this order of magnitude or better.

If the tags are moving relative to each
15 other, range calculation errors can result due to Doppler shift and to actual range changes. Velocities of less than 15 fps (feet per second) produce only insignificant Doppler shifts in the baseline method which uses a measurement interval of milliseconds. Errors due to
20 velocity are measurable above 15 fps and become a concern above 500 fps, where they are addressed using an alternate method as described below.

(b) Alternate Methods

Because the baseline method depends on the
25 originator sending pulses at a uniform rate and energy consumption requirements dictate a limited measurement duration, there is a maximum range at which the baseline method can function. Beyond that range, the range measurements have a range ambiguity. A range ambiguity
30 results when the originator does not wait for a first pulse to return before sending a second pulse. In such a case, the originator cannot determine whether a returning pulse corresponds to the first pulse sent or the second pulse sent. With the delays of the transponder and
35 originator being quite large compared to the delay

- 84 -

associated with the range, the transponder and originator delays can readily be set so that they are consistent with the short ranges expected in an asset management application. Other applications of the baseline method
5 require resolution of range ambiguities.

One variation of the baseline method that avoids the range ambiguities is illustrated in Fig. 40. In this variation, the originator transmits a pulse and determines an initial clock count for return of the pulse
10 (step 660), establishing a reference. Because a clock count always corresponds to one of two possible time periods, for example, "td" or "td+ts" in the discussion above, the initial clock count is either a low count value, for example, corresponding to "td," or a high
15 count value, for example, corresponding to td+ts. The originator then transmits another pulse and determines another count (step 662). If the count does not equal the high count or the low count (step 664), the interval between pulse transmissions is incremented (for a count
20 lower than the low count) or decremented (for a count higher than the high count) by one originator clock period (step 666). These increments and decrements are made continuously, adapting the range measurement to longer ranges and retaining the one-clock-period
25 difference between the originator and transponder clocks. As long as the originator keeps track of when each pulse is transmitted, range ambiguities are avoided, because the originator can distinguish among returning pulses by calculating when to expect each pulse's return.

30 As discussed above, the baseline method sets the clock frequencies such that the originator clock has one more or one less clock period during the entire measurement period. For the approximate magnitudes discussed above, this setting results in clock
35 frequencies which are different by about 50 ppm.

- 85 -

Alternate clock choices are used effectively to increase resolution. The primary requirements are (1) that the difference in the clocks' periods causes aliasing, and (2) that the clocks uniformly sample the measurement window, that is, the "ts" window, during the measurement interval. To ensure aliasing, the number of clock periods used, "m," is chosen such that "m" does not share any multiplicative factors with "N" in the measurement interval. Thus the relationship between "ts" and "tt" becomes:

$$t_s = t_t \frac{((N) (\text{spulses}) - m)}{((N) (\text{spulses}))} \quad (22)$$

where "N" and "m" share no common factors. The lack of common factors can be ensured by making "N" a prime number such as 101.

When velocity is a potential concern, as discussed above, two range measurements are made back to back. By dividing the difference of the two measured ranges by the time interval between the two measurements, the instantaneous velocity, including direction, is computed and used to correct the range measurement. If the transponder is moving closer to the originator, the range decreases during the measurement period. A correction equal to the magnitude of the range decrease, (velocity) (N) (spulses) (ts), is added to the range calculated. Conversely, if the range is increasing, the correction is subtracted.

As a result of limits due to the secondary errors discussed above, the duration of the measurement interval cannot be arbitrarily large in the baseline method. This is addressed by yet another variation of the baseline method. Referring again to Fig. 35, time as represented therein proceeds from left to right and the right end of each horizontal timeline represents an earlier time than the left end of the next horizontal

- 86 -

timeline below. If a second round of measurements were to be initiated "spulses" clock periods after the last pulse of the previous round, the results would repeat exactly, except for slight perturbations due to clock frequency errors. However, if the second measurement round is initiated $(n)(\text{spulses})/2$ clock periods later (where "n" is a positive odd integer), the effect is to shift all of the transponder detections in the "ts" window by $d/2$. This $d/2$ shift for an n-odd multiple of a half "spulses" period follows from the fact that the increase in ambiguity "x" for a full "spulses" period is "d." The sawtooth error profile shifts by $(c)(d)/2$ for the same reason.

After the two range measurements are averaged, the error profile of the result is still a sawtooth, but the error profile's period and amplitude are only half of the period and amplitude of each individual error profile. With reference to Fig. 38, halving of the period and amplitude occurs because the second round of measurements doubles the number of nominal range points. The doubling occurs because averaging the first and second rounds effectively adds a nominal range point halfway between each set of nominal range points of the first round. For example, an actual range which corresponds to a particular first round nominal range point and a different second round nominal range point is given by the average to correspond to an in-between point. The in-between point is a more accurate estimate of the actual range than was the first-round measurement. Thus, resolution is enhanced by averaging the first-round measurement and the second-round measurement.

In the full alternate method of this averaging technique, this process is repeated with four

- 87 -

or more measurement rounds. The number of rounds is divisible evenly into "spulses." In effect, this method is a binary search on the error space of the range measurement, reducing the ambiguity and maximum error for each measurement made. The reductions occur for the same reasons given above with respect to Fig. 38, namely, that additional measurements, spaced properly, effectively add additional nominal range points, thus providing an enhanced resolution.

10 In another variation, a binary search is undertaken in another fashion. In the baseline method, the delay "d" is a constant. However, in this variation, the delay d is varied to produce a binary search of the range space after identification of the measurement in
15 which a clock count transition from a lower value to a higher value occurs. For example, referring to Fig. 35, such a transition occurs in measurement 8, because in measurement 8 the count changes from "td" (measurements 1-7) to "td+ts" (measurements 8-10). The process of
20 halving the range error is continued in order to home in on the actual transition time and, in turn, the actual range. The successive delay in the measurements is determined by $d = (tt-ts)(spulses)$. This delay is increased by increasing pulse transmission intervals
25 "spulses" or by increasing the difference between the two clock periods.

 Because all tags can operate at two frequencies, corresponding to clock periods "ts" and "tt" as described above, an alternative strategy which
30 accomplishes the same homing-in result involves reversing the frequencies of the two tags as needed. With this strategy, described in detail below, if a timing difference of "-d" is needed, the originator tag commands the transponder tag to adopt the frequency corresponding
35 to clock period "ts." At the same time, the originating

- 88 -

tag makes the determination to set its frequency to correspond to clock period "tt."

This alternate method is now described in detail in an example which uses "N" = 8 (Figs. 41 and 5 42). At rest, the originator and the transponder use the same clock period "tt." Before the time 4100 when the first measurement is taken, the originator signals the transponder that a measurement is about to begin, then drops its clock period to "ts" and sends a pulse (step 10 676). To establish a reference clock count, the originator counts clock periods, which also reflect the transponder's delay "tdelay," until the pulse is returned 4102 (step 678). At this point, the originator starts to perform measurements at intervals corresponding to 15 "spulses" (step 680). Measurements are performed until the first to occur of:

- (1) A transition from an initial "low" count to a subsequent "high" count is identified 4104 (step 682). As described below, the 20 originator tag determines how to change the delays to focus on the identified period and commands the transponder tag accordingly.
- (2) All eight measurements have been taken and the counts are all the same, meaning that 25 there has been no transition from the low count to a high count (step 684). In that event, the originator resets its clock period for an amount of time corresponding to spulses/2 (step 686) and the preceding 30 sequence is repeated.

When the counts change from "low" to "high" as shown for the sixth measurement 4104 in Fig. 41, the

- 89 -

measurement method has detected that the range delay is such that an accumulation of 5 delays "d" is sufficient to "push" the return of the pulse into the "high" clock count period. Thus measurement 6 of Fig. 41 is similar to measurement 8 of Fig. 35. At this measurement, the measurement method checks, as follows, to determine more precisely the accumulation that is sufficient to have the same effect. Clearly, an accumulation of (4)(d) is insufficient because the fifth measurement did not produce the "high" count.

To test the effect of an accumulated delay of only (4.5)(d), the measurement method subtracts (0.5)(d) from the accumulated delay by switching clock periods in the following way. The originator stays in step with the transponder by raising its clock period to "tt" for one "spulses" period 4106 and then commands the transponder to lower its clock period to "ts" for half a "spulses" period 4108 (step 688). Because the originator and the transponder have different clock periods for only half an "spulses" period, only (0.5)(d) is subtracted from the accumulated delay. Another measurement and clock count are acquired and tested (step 690).

At this point, a clock count which remains high indicates that an accumulated delay of (4.5)(d) is sufficient to push the return into the high clock count period. In that case, an accumulated delay of (4.25)(d) is similarly tested, by lowering the transponder's clock period to "ts" for a quarter of an "spulses" period (step 688). On the other hand, a clock count which drops to the low count during the test of an accumulated delay of (4.5)(d) indicates that a measurement corresponding to an accumulated delay of (4.75)(d) should be taken (step 692) and tested (step 694). In this loop (steps 688, 690, 692, and 694), successively smaller portions of delay "d" are used to determine, to a successively higher

- 90 -

resolution, how much accumulated delay is sufficient to effect a count transition. The range delay follows directly from subtracting "tdelay" and the accumulated delay from the low count.

5 These steps thus, as described, implement a binary search. (Other search patterns could be used.) A range delay determined at this resolution still has the biases mentioned for the baseline method. And in a manner similar to that used in the baseline method, these
10 biases can be taken into account as the measured range delay is multiplied by the speed of light to produce the range.

 Finally, the same techniques of using two clocks also work using a single dual-clock tag with a
15 clockless passive or active second tag or reference. In a basic version of this single dual-clock tag variation, with reference to Fig. 43, one clock, having a period "tt," is used solely to cause the pulse to be transmitted 4200. The other clock, having a period "ts," begins
20 clocking the pulse when it is transmitted 4202 and then acts as the measure of time (in clock periods) to determine when it is received 4204. The number of "ts" receipt clocks in the interval between pulse transmission and detection of receipt of the returning pulse is the
25 basis of the range measurement.

 In this instance, "td" is a multiple of "ts" such that, as shown for transmission 4 in Fig. 43, "td" is an accurate measurement of the round-trip time when there are no leading delays, which corresponds to the
30 time (td-ts) in transmission 1 in the baseline method illustrated in Fig. 35. Thus, for transmission 4 here (in Fig. 42), the average range delay is simply calculated as half of (td-tdelay). That the round trip is measured as "td" in transmission 4 is indicated by the
35 fact that the return 4208 of the pulse is detected within

- 91 -

two (ignoring the hash marks in this illustration) receipt clocks of the pulse's transmission 4210. Transmissions 5-8, on the other hand, produce round-trip times of (td+ts), because there are three such receipt
 5 clocks between detection of transmission and detection of return. Thus, similar to the calculation of the range in the baseline method, every round-trip time of (td+ts) adds (ts/N) in the calculation of the range delay. As stated in general form, when "N" measurements have
 10 occurred and in "n" of the "N" measurements, the round-trip time corresponds to (td+ts), the measurement is given as

$$range_{aver} = \frac{c(td - t_{delay} + \frac{(n)(ts)}{N})}{2} \quad (23)$$

$$= c \left(\frac{(ts)(Totalclockcount)}{2N} \right) \quad (24)$$

The primary error of this basic dual-clock tag variation has a maximum magnitude of (c)(d)/2, which
 15 is the same as that of the baseline method. This is true because both the early clock (detecting pulse transmission) and the late clock (detecting receipt) have an initial ambiguity similar to the initial ambiguity "x" of the baseline method (Fig. 35). Thus each of these
 20 initial ambiguities has a maximum magnitude of "d" with respect to the transmit clock. Each maximum magnitude "d" is then halved and multiplied by the speed of light "c" in the range calculation to produce a range calculation error having a maximum magnitude of (c)(d)/2.
 25 The sum of these two range calculation errors has a minimum value of zero and a maximum value of (c)(d), which correspond to a minimum possible actual range and a

- 92 -

maximum possible actual range. Therefore, the primary error is $(c)(d)/2$, because the average range, which is calculated according to the equation above and which is the average of the minimum possible actual range and the maximum possible actual range, is always within $(c)(d)/2$ of the actual range.

The range calculation provides for an instance where the clockless second tag or reference adds a delay "tdelay." If the reference is a simple reflector, "tdelay" is set to zero. The binary search methods described above apply to this variation as well.

In further variations of the methods described above, each "pulse" described herein is actually a sophisticated symbol. The symbol is a complex waveform which is, for example, a combination of pulses or a spread-spectrum signal. A different combination of pulses or combination of frequencies is used for each symbol, which amounts to a unique signature for each symbol. A system using symbols with unique signatures overcomes some range ambiguity problems, because each unique signature is used for making an identification of a returning signal. In addition, use of the symbol provides for finer grain time detection and overcomes hardware-based secondary errors such as threshold noise.

25 (c) Energy Consumption

With the parameters specified above, the originator transmits during the measurement interval about 0.1% of the time. Both the originator and the transponder are left on for the entire measurement interval. If a measurement is made once per second on average, energy consumption is about 0.2-0.3% of full-time on-energy consumption.

However, after the first measurement interval is completed and the range is determined to a low resolution, both tags know at what time the next pulse

- 93 -

transmission or reception is to occur. Therefore, alternatively, the tags program themselves to "power-up" at that time, reducing energy consumption by a factor of about 100.

5 B. The verify range subfunction

Referring to Fig. 44, the verify range subfunction determines whether each range received from the intertag distance subfunction is in error, and creates a list 800 of ranges associated with an error. List 800 stores the
10 identification of the pair of tags associated with the range, the value of the range, the estimated error in the range (σ) as provided by the intertag distance subfunction, and the number of errors associated with the range.

15 Referring to Figs. 45 and 46, the subfunction begins by setting N equal to 3, and selects a group of N tags 808, 810, 812 (step 802). The subfunction tests whether the ranges a, b and c between the three tags form a triangle (steps 813, 814, 815), that is, whether within
20 expected measurement error, the sum of two ranges (a+b) exceeds the third side (c) or the magnitude of the difference (a-b) exceeds c. This condition is expressed as follows:

$$|a-b| - 2\sigma_c \leq c \leq a+b + 2\sigma_c \quad (25)$$

where σ_c is the standard deviation of an error metric,
25 defined by ϵ_1 and ϵ_2 , as follows:

$$\begin{aligned} \epsilon_1 &= c - |a+b| \\ \epsilon_2 &= c - |a-b| \end{aligned} \quad (26)$$

$$\sigma_c^2 \equiv \sigma_{\epsilon_1}^2 = \sigma_{\epsilon_2}^2 = \sigma_a^2 + \sigma_b^2 + \sigma_c^2$$

where σ_a , σ_b and σ_c are the measurement errors of a, b and c, respectively. ϵ_1 takes on values that are less than or equal to zero, and ϵ_2 takes on values that are not less than zero, when σ_a , σ_b and σ_c are negligible. Other

- 94 -

definitions of the error metric are possible; the definition provided in Equation (25) makes optimal use of the error in the ranges a, b and c.

Equation (26) assumes the variability in the values of a, b and c, due to measurement error is accurately modelled by a Gaussian probability distribution, centered at the measured value of a, b and c that is being tested, and with a standard deviation σ_a , σ_b , σ_c which are all statistically independent of one another. Equation (27) is derived by bounding ϵ by its standard deviation (although the multiplier does not necessarily have to be two):

$$\begin{aligned} \epsilon_1 &\geq -2\sigma_\epsilon \\ \epsilon_2 &\leq 2\sigma_\epsilon \end{aligned} \quad (27)$$

An error is detected when c does not fall within the bound given in Equation (25). When this occurs, the subfunction increments the number of errors for ranges a, b and c in list 800 (step 816). Referring to Fig. 47, the subfunction also creates an error log 820, and appends an entry to the log for each triplet that gets flagged as an error in list 820 (step 822). These entries include all the ranges associated with a particular error.

The subfunction then selects a next group of three tags (step 823), and returns to step 813. The subfunction continues the process until each triplet of tags in the network has been analyzed. If no error is detected in step 815, the subfunction proceeds to step 823 and analyzes another group of three tags.

After all N-tuples have been checked, the subfunction finds the range R with the highest number of error flags in list 800 (step 824). For each triplet in the log 820 that includes the range R (step 825), the subfunction decrements the number of errors in list 800 for each range in the triplet (step 826). The

- 95 -

subfunction then removes the triplet from the log 820 (step 827) and chooses a next group of three tags that includes the two tags forming the range R (step 828).

After the subfunction has checked every triplet in
 5 the log that includes range R, the subfunction checks whether any errors remain in list 820 (step 829). If errors remain, the subfunction returns to step 824 and repeats the process. If no errors remain, the subfunction sets $N = 4$ (step 830) and returns to step
 10 802.

As shown in Fig. 48, the subfunction next selects a group of four tags (A, B, C and D), and retrieves the ranges (a,b,c,d,e,f) between each of the six pairs of tags (step 813). Steps 814 and 815 involve the following
 15 analysis.

The range verification subfunction models the tags as forming two triangles (ABC) and (ABD) with a common side (AB). The subfunction chooses the longest range (e) as the common side (AB). For a given set of ranges,
 20 there is an ambiguity in the location of D. The possible locations of tag D are indicated by D_1 and D_2 . The distance between C and D_1 is labelled f_{\min} ; the distance between C and D_2 is labelled f_{\max} .

The subfunction calculates f_{\min} and f_{\max} from the
 25 ranges a, b, c, d and e. The equations necessary for this calculation are given in Appendix I. The subfunction then checks whether the four tags form a valid tetrahedron by verifying whether the range f is bounded as follows:

$$f_{\min} - 2\sqrt{\sigma_{f_{\min}}^2 + \sigma_f^2} < f < f_{\max} + 2\sqrt{\sigma_{f_{\max}}^2 + \sigma_f^2} \quad (28)$$

30 where σ_f is the error in f, $\sigma_{f_{\min}}$ and $\sigma_{f_{\max}}$ are Gaussian approximations for the variances of f_{\min} and f_{\max} , respectively (step 854). Equation (28) is derived by bounding f such that it does not vary from its mean by

- 96 -

more than two standard deviations. In other embodiments, f is bounded by a different number of standard deviations.

Equation (28) also assumes that each range measurement (a,b,c,d,e,f) has a Gaussian profile, and the ranges are uncorrelated with one another. A Gaussian approximation for f_{\min} and f_{\max} is derived by linearizing the expression relating f_{\min} and f_{\max} to the other ten ranges (a,b,c,d,e), as shown in Appendix I below.

10 Referring to Fig. 49, if f does not fall in the bound given by Equation (28), the subfunction increments the number of errors in the stores associated with the ranges a through f in list 800, and creates an error log 820' analogous to error log 820. The subfunction appends
15 an error flag to each range in the group of four tags in log 820'. The subfunction then performs the remaining steps in Fig. 45 using list 800 and log 820'.

When the subfunction has removed all the errors from list 800 (step 829), the subfunction sets N to 5
20 (step 830) and returns to step 802. The subfunction then selects a grouping of five tags and retrieves the group of ten ranges (a,b,c,d,e,f,g,h,i,j) between each pair of tags in the grouping (step 813). Steps 814 and 815 involve the following.

25 The subfunction models the five tags as two tetrahedrons sharing a common triangle. The subfunction chooses the triangle with the largest area as the common triangle.

Assuming the other nine ranges are correct, j is
30 geometrically constrained to have one of two values (j_1 or j_2). When at least one of the ranges is in error, j is bounded by:

- 97 -

$$\{|j-j_1| \leq 2\sqrt{\sigma_{j_1}^2 + \sigma_j^2}\} \cap \{|j-j_2| \leq 2\sqrt{\sigma_{j_2}^2 + \sigma_j^2}\} \quad (29)$$

where σ_j is the measurement error in j , and σ_{j_1} and σ_{j_2} are the standard deviations of j_1 and j_2 when j_1 and j_2 are approximated by Gaussian functions. Equation (29) is derived by bounding j such that it does not vary from its mean by more than two standard deviations. In other
 5 embodiments, j is bounded by a different number of standard deviations. Equations for deriving j_{\min} and j_{\max} and the variances are given in Appendix II.

The subfunction verifies whether j falls within
 10 the bounds described by Equation (29) (steps 814, 815). Referring to Fig. 50, if the condition is not satisfied, the subfunction increments the error in the stores associated with a through j in list 800, and creates an error log 820" which is analogous to error logs 820,
 15 820'. The subfunction appends an error flag to the log for the pairs of tags in the grouping of five tags, and repeats the remaining steps in Fig. 45.

When the process returns to step 830 this time, the tag proceeds to the refine range subfunction step
 20 830a, described below.

C. The refine range subfunction

The refine range subfunction estimates a more accurate value for each range. Even valid range measurements can be improved with this technique.

25 The subfunction first selects a range (f) to be estimated, and finds groups of five tags that define pairs of tetrahedrons sharing the range in a common triangle. The processor calculates a value for the common range based on the nine measured ranges
 30 ($a, b, c, d, e, g, h, i, j$) in each group of five tags. This yields a group of estimates (m_1, m_2, \dots, m_N) for the common range (assuming N tetrahedron pairs are used).

- 98 -

As shown in Appendix III, each estimate is linearized to x_1, x_2, \dots, x_N , and modelled as a Gaussian probability function $f_1(x), f_2(x), \dots, f_N(x)$ with a standard deviation equal to the measurement error of the estimate $(\sigma_1, \sigma_2, \dots, \sigma_n)$ (calculated, as is well known in the art, from the standard deviations of the measured ranges a, b, \dots, j for each tetrahedron pair) and a mean centered at the estimate itself (m_1, m_2, \dots, m_N) .

The best estimate x_b is modelled by a Gaussian probability function $f_b(x)$ with a mean equal to the estimate x_b and a standard deviation σ_{xb} . The best estimate is derived by finding x_b that has a maximum likelihood, given the Gaussians $f_1(x), f_2(x), \dots, f_N(x)$. This yields:

$$x_b = \frac{\sum_{i=1}^N \sum_{j=1}^N (x_i + x_j) q_{ij}}{2 \sum_{i=1}^N \sum_{j=1}^N q_{ij}} \quad (30)$$

where q_{ij} is the (i, j) th element of an inverse covariance matrix of the set of estimates x_1, x_2, \dots, x_N .

Instructions for computing the covariance matrix are given in Appendix III.

The variance of the estimate x_b is given by:

$$\sigma_{x_b}^2 = \frac{1}{\sum_{i=1}^N \sum_{j=1}^N q_{ij}} \quad (31)$$

After performing the calculation given in Equation (30) the processor assigns the best estimate x_b to the range f . The processor then selects another range, and repeats the process until each range has been refined.

D. Compute locations

Once refined range measurements are obtained for each pair of tags, the compute locations subfunction sets up a coordinate system for defining the position of each

- 99 -

tag. The coordinate system is defined by four stationary tags, as follows.

Referring to Figs. 51 and 52, a pair of tags 16_1 , 16_2 , that are farthest apart defines the x-axis of the coordinate system (step 901). One of the tags 16_1 is chosen as the origin, while the other tag lies on the positive x-axis (step 902). Choosing the tags furthest apart renders the coordinate system less sensitive to measurement error in the ranges between the tags. This coordinate system is transformed into an absolute coordinate system in section E below.

Next, the processor finds the tag 16_3 , that is furthest from the x-axis. This tag along with the x-axis defines the xy plane (step 904). The third tag is assigned a positive y coordinate. The known ranges between the three tags are used to define the coordinates of the third tag with respect to the origin, using conventional geometry and trigonometry.

The processor then chooses the tag 16_4 , furthest from the xy plane, and assigns a positive z value to the tag (step 906). This completes the definition of the coordinate system.

The x and y coordinates and the magnitude of the z coordinates of the remaining tags are defined by triangulation with the first three tags. The fourth tag is used to define the sign of the z coordinate (step 908).

E. Absolute Coordinate System

The arbitrary coordinate system of Fig. 51 is transformed into an absolute coordinate system, if the location of four stationary tags (which know their location) was supplied by the intertag distance subfunction. It is important that none of the tags are co-located (that is, located at the same coordinates) or coplanar.

- 100 -

If the locations of four stationary tags are not known, the tag checks whether the intertag distance subfunction supplied a list of moving tags. If so, the subfunction retrieves the list of moving tags. The list
5 includes the locations of the tags and the error in the location measurement at given times, along with the distances between the tag and the moving tags at those times.

The subfunction analyzes this information along
10 with the locations of the stationary reference tags to determine whether one of the following conditions holds for a tag. If any one of the conditions is satisfied, the tag determines its instantaneous position at each of the times relative to the reference objects with a
15 maximum likelihood method analogous to that used by the refine range subfunction. The method employs all the information in the list to derive an estimate of the relative position of the tag and an error for the estimate.

20 The tag then decides whether the tag itself is moving, a reference tag is moving, or both. If the tag itself is moving, the tag determines its instantaneous position, and velocity using the relative position information. If the tag is stationary, the tag converts
25 its relative position to an absolute position.

The first condition is that the reference tags must be moving in three dimensional space relative to the tag.

If the tag and reference tags are coplanar but not
30 co-linear, the reference tags must be moving in two dimensions relative to the tag. If the tag and reference tags are co-linear, the reference tags must be moving in one dimension relative to the tag.

As additional reference positions and errors are
35 accumulated by the intertag distance subfunction, the tag

- 101 -

refines its position estimate using the same method. In addition, as tag locations are identified and changed, this information is passed to the determine topology subfunction.

5 XI. Detailed description of network formation unit

A. Determine topology subfunction

The determine topology function designates a set of neighbors for each tag by means of Delaunay triangulation, described in detail in "Spatial
10 Tessellations: Concepts and Applications of Voronoi Diagrams," A. Okabe et al, the contents of which are incorporated herein, in its entirety, by reference.

Referring to Fig. 53, the neighbors of a tag 16₁ are found by first locating the set of points that are
15 closer (in Euclidean distance) to tag 16₁ than to any other tag. Each set of points is bounded by an "ordinary Voronoi polygon" 2000. The neighbors of each tag include the tags in polygons adjacent to polygon 2000. Thus, the neighbors of tag 16₁ consist of tags 16₂, ..., 16₈.

20 Referring to Fig. 54, a preliminary communication link 2010 is defined between a tag and each neighboring tag. The set of links between tags forms a Delaunay triangulation, and defines the topology of the set of tags.

25 The Delaunay triangulation is defined incrementally by causing each tag to find its neighbors. If the tags are all activated at the same time, the netmaster (the tag with the highest ID number) finds its neighbors first, by using the algorithm described below.
30 After the netmaster has completed the algorithm, it selects the tag having the second highest ID number, and commands the tag to perform the algorithm to find its neighbors. The netmaster continues selecting tags in decreasing order of ID number until all tags have been
35 selected.

- 102 -

When a new tag is added to a network, it communicates with other tags to determine whether the tags have already performed the process described above, that is, whether a Delaunay triangulation for the tags has been determined. If it has, links must be added to the triangulation to connect the new tag to its neighbors. This is accomplished by having the new tag perform the algorithm described below. If no Delaunay triangulation exists, the new tag waits until ordered by the netmaster to perform the algorithm.

The new tag (or the tag designated by the netmaster) finds and maintains a list of its neighbors by performing the following algorithm. Each step of the algorithm is executed by the determine topology subfunction in the new tag.

The new tag begins by receiving its own coordinates and then seeks information locating other tags in its vicinity from the compute locations subfunction. The determine topology subfunction ignores the z coordinates in the location information and thus finds the topology of the tags using a two-dimensional projection of their locations. The determine topology subfunction next selects a random tag from the coordinate information which it defines as the "current tag" (current(n)) where n is the index of the iteration and queries the current tag regarding its neighbors. In other embodiments, the tag may communicate with a sentinel (defined in Section IX.C(c), above) that causes the tag to choose the sentinel as the current tag. Any other tag may be chosen as the current tag.

Referring to Fig. 55, the new tag calculates its distance from the current tag and its distance from each neighbor of the current tag, using the coordinate information (step 2014). The new tag analyzes the distances to determine whether the new tag is closest to

- 103 -

the current tag (steps 2016, 2018). If so, the new tag stores the identification number of the current tag in its list of neighbors (step 2020) and proceeds to step 2022 in Fig. 56. If not, the new tag chooses the
5 neighbor of the current tag that is nearest the new tag as the current tag (step 2022) and returns to step 2014.

When the new tag arrives at step 2022, the current tag has been defined as either the last current tag or the neighbor of the last current tag that is
10 closest to the new tag. In step 2024 (Fig. 56), the new tag defines two of the neighbors of the current tag as the "next tag" and the "previous tag."

The next tag of a particular tag A is defined relative to a second tag B. The next tag is the neighbor
15 of tag A which forms the smallest positive angle c with respect to tag B, with the tag A being placed at the vertex of the angle. For example, as shown in Fig. 57, tag A is the current tag 16_1 , and the tag B is the new tag 16_2 . If the neighbors of the current tag include tags
20 16_3 , 16_4 , and 16_5 , then the next tag of the current tag relative to the new tag of the current tag is tag 16_3 . Similarly, the next tag relative to tag 16_3 is tag 16_4 .

The previous tag of a current tag relative to a second tag is defined as follows. The previous tag is
25 the neighbor of the current tag which forms the smallest negative angle d with respect to the second tag, with current tag placed at the vertex of the angle. In the example shown, the previous tag for the current tag 16_1 with respect to the new tag 16_2 is the tag 16_4 .
30 Similarly, tag 16_3 is the previous tag for the current tag relative to tag 16_4 .

The new tag defines the next tag ($\text{next}(n)$) and the previous tag ($\text{previous}(n)$) of the current tag relative to the new tag (step 2024). The new tag then checks whether
35 the angle between the previous tag and the next tag (with

- 104 -

the current tag at the vertex of the angle) is less than 180 degrees (step 2026). For example, in Fig. 57, the new tag checks whether the angle d is less than 180 degrees. In the same step, the new tag verifies that the
5 previous and next tags are neighbors of one another (that is, they are connected in the Delaunay triangulation). The new tag accomplishes this by querying the previous tag and the next tag regarding their neighbors.

If both conditions of step 2026 are satisfied, the
10 new tag checks whether the new tag is located inside a circle circumscribing the current, next and previous tags (step 2028). For example, as shown in Fig. 58, circle 2029 circumscribes tags 16_1 , 16_3 , 16_4 , and contains the new tag 16_2 . If the condition of step 2028 is satisfied, a
15 new previous tag (previous $(n+1)$) is defined as the next tag (next (n)), and a new next tag (next $(n+1)$) is defined as the tag having the smallest positive angle with respect to the next tag (next (n)). The new tag then increments the value of the iteration ($n=n+1$) and
20 returns to step 2026.

If the condition in step 2026 is not satisfied, the new tag checks whether the new tag is between the previous tag and the next tag with respect to the current tag (step 2034). This means that the previous and next
25 tags are previous and next tags relative to the new tag. For example, in Fig. 57, the new tag 16_2 is located between the previous tag 16_5 and the next tag 16_3 . If the condition in step 2034 is satisfied, the new tag adds the previous tag to its list of neighbors, and informs
30 the previous tag. The previous tag then adds the new tag to its list of neighbors (step 2036). Similarly, if the condition in step 2028 is not satisfied, the new tag proceeds to step 2036.

After step 2036, the new tag checks whether the
35 previous tag is identical to the closest tag found in

- 105 -

step 2018 (step 2037). If so, the algorithm skips to step 2038 (Fig. 60). If not, the new tag checks whether the angle of the new tag to the previous tag (with the current tag at the vertex of the angle) is greater than
5 or equal to 180 degrees (step 2040). For example, in Fig. 57, the new tag compares the angle d to 180 degrees. If the condition is not satisfied, the current tag (current(n+1)) is redefined as the previous tag (previous(n)) (step 2042), the index n of the iteration
10 is incremented by one (step 2044), and the new tag goes to step 2024. If the condition in step 2040 is satisfied, the algorithm proceeds to step 2022' (Fig. 59).

Referring to Fig. 59, in step 2022', the algorithm
15 begins analyzing the previous and next tags in an opposite order. Step 2022' is identical to step 2022 in Fig. 56. The algorithm next repeats steps 2024 to 2044, with the following changes. In step 2030' (which replaces step 2030), a new next tag (next (n+1)) is
20 defined as the previous tag (previous(n)). A new previous tag (previous(n+1)) is defined as the tag forming the smallest negative angle (with the current tag at the vertex of the angle) relative to the previous tag (previous(n)). In step 2036' (which replaces step
25 2036), the next tag (next (n)) is added to the list of neighbors of the new tag. The new tag then informs the next tag, which adds the new tag to its list of neighbors. Steps 2037 and 2040 are replaced with step
30 2037', in which the new tag checks whether the next tag is the closest tag, or the angle between the new tag and the next tag (with the current tag at the vertex of the angle) is greater than or equal to 180 degrees. If the condition is not satisfied, the algorithm proceeds to step 2042. If the condition is satisfied, the algorithm
35 proceeds to step 2038 (Fig. 60).

- 106 -

Referring to Fig. 60, in step 2038, the new tag finds a pair of its neighbors (A, B) that are neighbors of each other, but are not consecutive with respect to the new tag. This means that tag A is not the neighbor
5 of the new tag that forms the smallest positive angle with tag B, with the new tag placed at the vertex of the angle. Similarly, tag A is not the neighbor of the new tag that forms the smallest negative angle with tag B, with the new tag placed at the vertex of the angle.
10 Fig. 61 shows a situation where the condition is satisfied, with tags A, B and C as neighbors of new tag 16₂.

In step 2038, the new tag also checks that tags A and B form an angle smaller than 180 degrees (with the
15 new tag placed at the vertex of the angle). In the example shown in Fig. 61, the new tag compares angle ϵ to 180 degrees. The algorithm checks every possible pair of neighbors, and if no pair of neighbors satisfy the conditions in step 2038, the algorithm ends. For each
20 pair satisfying the conditions of step 2038, the new tag checks whether one of the following two conditions are satisfied: (1) the next neighbor of tag A relative to tag B, and the previous neighbor of tag B relative to tag A are the same tag (tag C), and (2) the next neighbor of
25 tag B relative to tag A and the previous neighbor of tag A relative to tag B are the same tag (tag C) (step 2050). If either condition is satisfied, the new tag checks whether tag C is not the new tag and the new tag is included in a circle circumscribing tags A, B and C (step
30 2052). If the condition is satisfied, the tag disconnects tags A and B, that is, the new tag informs tags A and B that they are not neighbors of one another, and tags A and B modify their lists of neighbors accordingly (step 2054). The new tag then returns to
35 step 2038.

- 107 -

If the condition in step 2052 is not satisfied, the algorithm returns to step 2038. If the condition in step 2050 is not satisfied, the new tag checks whether the angle of the next tag of tag A relative to tag B
5 (with the vertex of the angle at tag A) or whether the angle of the next tag of tag B relative to tag A (with tag B at the vertex of the angle) is less than 180 degrees (step 2056). If so, the new tag proceeds to step 2054; if not, the algorithm returns to step 2038.

10 When a tag is removed or deleted, the netmaster directs updates of the topology by having the following algorithm performed by the neighbors of the removed tag. Each step of the algorithm is performed by the determine topology subfunction in the netmaster.

15 Referring to Fig. 62, the netmaster identifies a current tag as a neighbor of the removed tag (step 2100) and directs it to execute the following. The next tag (next(n)) is defined as the neighbor of the removed tag that is the next tag relative to the current tag with
20 respect to the deleted tag. The previous tag (previous(n)) is defined as the neighbor of the removed tag that is the previous tag relative to the current tag with respect to the deleted tag (step 2102). The tag checks whether the next tag is the same as the previous
25 tag and whether the previous and next tags are neighbors of one another (step 2104). If either or both conditions are satisfied, the algorithm causes each neighbor of the deleted tag to remove the deleted from its list of neighbors. The deleted tag is thus disconnected from all
30 its neighbors, and the process ends. If the conditions of step 2104 are not satisfied, the netmaster checks whether the current tag is connected to the next and the previous tags, that is, the current, previous, and next tags are neighbors of one another (step 2106). If not, the tag
35 defines a new current tag (current(n+1)) as the next tag

- 108 -

(next(n)), increments the index n of the iteration (step 2108) and returns to step 2102. The removed tag completes the process by removing all the neighbors from its list.

5 If the condition in step 2106 is satisfied, the tag checks whether a neighbor of the deleted tag other than the current, next or previous tags is contained in a circle circumscribing the next, current and previous tags (step 2110). If the condition is satisfied, the
10 algorithm returns to step 2108. If not, the tag checks whether the deleted tag is contained in the circle (step 2112). If the condition is not satisfied, the algorithm returns to step 2108. If the condition is satisfied, the deleted tag is removed from the list of neighbors of the
15 current tag, and the next and previous tags are listed as neighbors of one another (step 2114). The algorithm then returns to step 2108.

 When a tag is moved from one location to another, the tags update the topology by performing the algorithm
20 illustrated in Fig. 62, in which the moved tag is defined as the deleted tag. The moved tag then finds its neighbors in its new location by performing the algorithm illustrated in Figs. 55, 56, 59 and 60.

B. Propagate messages

25 The propagate messages subfunction 564 (Fig. 25) defines a set of relay points for efficient communications between tags, and defines a hierarchy of networks.

 Each network within a level of the hierarchy is
30 defined by the physical constraints of the system. For example, in one type of system, the assets include items placed in a container. Each container is in turn stored on a pallet carried by a vehicle. An asset tag is attached to each item, pallet, container and vehicle in
35 the system. The set of tags on the items within a

- 109 -

container neighbor one another, and automatically form a network. Because the tag on the container is likely to be within the communication range of only a few tags on the items, the tag on the container is generally not
5 included in the network. Similarly, the tags on the packages, pallets, containers and vehicles form networks among themselves. Each network defines one level of the hierarchy.

A tag in each level chooses a sentinel or gateway
10 tag for communication with a tag in an adjacent level of the hierarchy. For example, the gateway tag on the container communicates with a tag on the pallet. A tag on an item may thus communicate with the tag on the pallet through the gateway in the container. Each tag is
15 thus able to communicate with any network in the hierarchy, without requiring a long range communication capability.

In some embodiments, tags may communicate with one another directly by bypassing an intermediate level in
20 the hierarchy. For example, a device may query a tag on an asset in a container directly, without accessing the gateway tag on the container.

In addition, each level in the hierarchy is not required to form a network, but may consist of tags
25 communicating independently with the next level of the hierarchy. In one important application, the hierarchy described above is implemented on a ship transporting containers storing items. The container creates a bill of lading for the items in the container and communicates
30 the bill to a device on the ship. The device then relays the bill to the outside world with cellular telephone, satellite relay or RF links. Alternatively, the container is loaded on a pallet that communicates with the device.

- 110 -

XII. Specific Applications

A. Hazardous waste transportation

(a) System design

Referring to Figs. 6, 63 and 64, one
5 application 10₄ of the system monitors trucks 12₆
transporting containers 12₁ of hazardous materials in a
trailer 12₆'. An asset tag 16₁ is attached to each
container. Alternatively, the containers are arranged on
pallets that are each provided with an asset tag, or a
10 group of pallets may be provided with a single tag. The
trailer also includes one or more sensors 18, for
example, a temperature sensor and a door sensor. A
helper tag 20₁ attached to an asset tag 16' is placed on
the inside of the trailer, and is in communication with
15 the sensors and the tags on the container. The helper
tag and asset tag form a gateway for communication
outside the trailer.

An interrogate and respond module 22 is
placed on the outside of the trailer, and is coupled to
20 the helper tag 20₁ on the inside of the trailer through a
hole in the trailer wall. Both the helper tag in the
interrogate and respond module and the helper tag on the
inside of the trailer are battery powered. In an
alternative embodiment, these tags are powered by the
25 trailer.

The tractor 12₆ of the truck is equipped with
an interrogate and respond module 22' and a monitoring
device, which includes a PDA 270 connected to a
communications controller 272 by a boot 20', asset tag 16
30 and holster 20," as shown in Fig. 20 above. The PDA has
a special panic button. The communications controller
has a cellular link 212' to an operations center 13. The
communications controller in the tractor is powered by
the vehicle, and provides power to the interrogate and
35 respond module 22' on the tractor. A GPS sensor 18' and

- 111 -

other sensors 18 (such as a temperature sensor) are also mounted in the tractor. The interrogate and respond module 22' on the tractor receives messages from the controller over an RS 232 serial link, and converts
5 the messages to RF signals which are relayed to the interrogate and respond module 22 on the outside of the trailer. Both modules 22, 22' are queried by interrogators 3000, which include police radar guns or similar devices typically installed at weigh stations,
10 toll booths and border crossings.

Referring to Fig. 65, the microprocessor software of each asset tag in this application is a simplified version of the microprocessor software described above. The microprocessor includes only a
15 self-assessment unit 370, a joining unit 372, and a network use unit 378. The self assessment unit includes the battery 400, active 402 and attachment 404 subfunctions. The joining function includes a simplified sentinel subfunction 562 that checks only for gateways
20 and seeks only the PDA (but not other asset tags). The network use function includes only the read and write data subfunction 596.

The asset tags 16, on the containers can be, typically, 30mm x 30mm x 15mm, and contain enough battery
25 power to last two years. The asset tag on the inside of the trailer can be about a 30 mm cube, while the asset tag in module 22 can be about 30 mm x 90 mm x 20mm. An asset tag including the GPS sensor can be about 30 mm x 90 mm x 30 mm in size.

30 (b) System operation

The tag on the inner trailer is powered on at all times, or, alternatively, it monitors its communication circuitry for messages before fully powering on. Upon receipt of a message, the tag checks
35 whether the message is destined for itself or another

- 112 -

component, in the manner described above. If the message is for itself, the tag responds to all the commands listed in Section V. The tag ignores and passes on all messages addressed solely to other tag ID's.

5 The helper tag on the inner trailer, when in an operations mode, sends a data retrieval message to the sensors in the trailer. The helper tags in the interrogate and respond modules 22, 22' perform the same functions as the helper tag in the inner trailer, but do
10 not respond to a command changing the time T during which the synchronization code is broadcast. This is because the helper tags communicate only with one another and the PDA. Because the helper tags in the modules and the PDA are powered on at all times, they will immediately
15 receive the synchronization code when broadcast by a helper tag. The helper tags therefore do not need to broadcast the synchronization code over time T, which is designed to allow a sleeping receiver enough time to wake up and receive the code.

20 Modules 22 and 22' and asset tags 16, on the containers 12, can communicate with modules on other vehicles. If an asset tag detects it has been detached, the tag relays the information to a module on its own vehicle or to the module on another vehicle if its own
25 module is disabled.

 In a typical scenario, a buyer contacts a shipper 3002 (Fig. 63) and requests certain virgin materials. The shipper communicates information regarding the buyer to the operations center 13. The
30 center in turn supplies the information to a carrier 3004 who transports the materials to the buyer.

 Alternatively, the shipper locates a buyer for its waste materials. The shipper then contacts the center in order to locate a carrier for transporting its
35 waste materials to the buyer.

- 113 -

In each of these scenarios, the shipper identifies to the operations center the materials to be shipped, their location and their destination. The shipper may also specify transshipment points, along with
5 dates and times for pickup and delivery at each point. The operations center 13 then organizes a route for one or more carriers based on this information, and relays the routes to the carriers' vehicles. Alternatively, the shipper may request that the system relay the shipper's
10 information to a carrier which organizes its own route. A shipper or buyer may contact the operations center to obtain the route for the materials, or to subsequently modify or delete his order.

Before shipment, the vehicle driver picks up
15 the containers of materials and attaches an asset tag to each container of materials. The user employs the PDA in the monitoring device to store the name of the shipper, the recipient and a description of the contents of the container on the tag, to activate the tag and to load
20 data into the tag. Each person who takes possession of the container also enters a signature on the PDA that is loaded onto the asset tag. This fulfills the shipper's legal requirement of obtaining the signature of each person who takes possession of the hazardous waste.

25 The operations center communicates route information to the PDA, as well as a password for the driver who will take possession of the materials. When the driver receives the materials and the PDA, the driver first enters his password into the PDA. After the PDA
30 verifies the password, it displays the route information to the driver.

Alternatively, the PDA does not store the password and route information. The operations center instead verifies this password and the route for the

- 114 -

materials is communicated back to the driver. Subsequent drivers repeat this process.

At each stop and periodically between stops, the PDA 270 communicates with each tag in the trailer and verifies that no containers are missing. The PDA also checks for extraneous tags. If discrepancies are found, the PDA notifies the monitoring device. If the user does not employ the device to notify the PDA that action will be taken, the PDA notifies the operations center of the discrepancy and the location of the vehicle. The center then takes action according to a prearranged plan developed with the carrier.

At the carrier's option, the PDA may access the modules to report the status of the assets (regardless of whether there is a discrepancy) periodically, randomly, at every stop, at particular locations, or after a fixed time from a stop. The PDA is loaded with route information that includes a series of expected times of arrival at specific locations and a tolerance for error. The PDA monitors the location of the vehicle through the GPS sensor and compares it to the plan. When a discrepancy beyond the tolerance for error is detected, the PDA reports the discrepancy to the center. The center responds to the discrepancy according to a prearranged plan with the carrier.

The operations center notifies PDA 100 of any changes in the route due to a request from the shipper, weather or traffic problems, or vehicle breakdowns.

At any time during the shipment, the shipper, buyer or a carrier may query the center regarding the location of the materials. The center verifies the requestor's legitimacy, and determines which vehicle is transporting the materials. The center then queries the PDA for its location and asks the PDA to verify that the materials are on board. Alternatively, if the PDA has

- 115 -

not reported any discrepancies with its planned route, the center deduces the location of the materials from the route.

When the materials reach their destination, the driver uses the PDA in the monitoring device to upload the signatures stored in each tag and to delete the signatures from the tag memory. He then deactivates each tag and removes it from the container. The device automatically relays the signatures to module 272, which in turn communicates them to the operations center. The center then supplies the signatures to the shipper. Modules 22 and 22' may be interrogated at any time with a police speed measuring radar. The modules broadcast the contents of the vehicle with synthesized voice on a police radio channel. The voice communication is followed by a data stream that may be monitored on special equipment.

If an incident occurs during shipment, the vehicle driver presses the panic button in the PDA and confirms by pressing it again. The PDA automatically forwards an emergency message along with the location of the vehicle to the operations center. The operations center contacts a first responder suitable for the location and materials, an appropriate police organization and the carrier. The center informs the responder and the police of the location of the vehicle, the types and quantities of materials being carried and telephone numbers of the carrier. Alternatively, only the carrier will be notified so that he may contact the vehicle driver and make appropriate arrangement with responders.

If the operator is incapacitated and cannot touch the panic button, policemen may interrogate the modules on the vehicle to determine its contents. The police may then take appropriate action. In addition,

- 116 -

the tags on the containers may be individually queried with a monitoring device.

(c) Other embodiments

Other relationships between shipper, buyers
5 and carriers are easily accommodated by the system. For example, several carriers may be coordinated to transport the materials from the shipper to the buyer.

In other embodiments, various components of the system described above are omitted. For example, if
10 a centralized service is not required, the operations center is omitted. In this scenario, all communications take place through a distributed system of monitoring devices.

If the center is included, the client
15 communication, inventory, scheduling and accounting software described above may be omitted. In addition, the Internet link, uninterruptable power supply, generator, and redundant backups may be removed from the center.

20 In other embodiments, a client is provided with a computer 14 (Fig. 1) for communicating with the operations center. The computer has inventory software for processing inventory information relayed by the center.

25 In some embodiments, the sensors in the tractor and the holster and boot in the communications controller are omitted. In addition, the communications controller communicates with an asset tag on the outer trailer using an asset tag on the tractor, but without
30 using an intervening interrogate respond module. The interrogate and respond module is thus replaced with a simple asset tag. The tag receives RF communication from the container tags, and relays the communications to the communications controller.

- 117 -

In other embodiments, the attachment function and the active subfunction are omitted from the container tag processors. In addition, the full network formation and location functions may be added. B. Warehousing

5 (a) Structure

Referring to Fig. 3, another application 10, of the system monitors equipment stored in a warehouse. In addition to use in a regular warehouse environment, the system is useful for monitoring valuable objects in
10 homes and museums, and for managing and locating objects in distributed locations, for example, in a hospital. Warehouse personnel attach an asset tag to each item in the warehouse, either at the entrance to the warehouse, at the origin of the item, or at some intermediate point.
15 The asset tags contain the full range of functions described in connection with Fig. 23.

Fork lift trucks and other conveyances in the warehouse are each equipped with a helper tag connected to both an asset tag and a GPS sensor. The helper tag is
20 powered by the battery on the conveyance, and provide this power to the external power port on the asset tag. The helper tag has an RS 232 port connected to the GPS sensor.

Pseudosatellites (also called "pseudolites")
25 located in the warehouse communicate with the GPS sensors and may communicate with GPS satellites outside the warehouse. One or more gateway tags 16, are disposed in the warehouse.

An inventory control computer 14 in
30 communication with the gateway maintains a list of assets in the warehouse. Warehouse personnel employ monitoring devices 15 for communicating with the asset tags, the gateway tags and GPS sensors.

The warehouse application does not include an
35 operations center. All information is collected by

- 118 -

warehouse personnel using monitoring devices, and is downloaded onto the local inventory computer. As a result, the monitoring devices do not include the communications controller described above, but still
5 employs a boot and an attached asset tag. The PDAs in the monitoring devices are recharged using a standard recharging cradle.

(b) Operation

The system is assembled by causing the
10 pseudolites to monitor GPS signals from the GPS satellite outside the warehouse. The pseudolites determine their own location based on the GPS signals, and communicate their location to the GPS sensor in the warehouse. The sensors on the conveyances monitor the GPS signals from
15 the pseudolites when in the warehouse, and monitor the GPS signals from the satellite directly when outside the warehouse. The helper tag communicates with the sensor over the RS-232 port, and converts the location information to an unmodulated IR signal. The helper tag
20 communicates the signal over its IR transmitter to the asset tag. The asset tag stores the location information.

Warehouse personnel add an asset to the system by choosing an asset tag for the asset and placing
25 the tag in the programming port of a monitoring device. The PDA in the device contains a list of assets to be tagged, each matched to an identification number and a set of data. The data includes the contents of the asset, the shipper, a list of waypoints and a
30 destination. If the asset is not being shipped, the destination is designated "inventory." The user chooses the number associated with the asset to be tagged, and causes the PDA to load the identification number and set of data onto the tag. The PDA then reads
35 the data stored on the tag, and verifies that it matches

- 119 -

the information downloaded. If the information stored on the tag is erroneous, the PDA repeats the download operation. If the information is correct, the PDA activates the tag and indicates to the user that the tag should be detached from the programming port and attached to the object.

If the particular asset is not listed on the PDA, the user enters the ID of the asset tag in the PDA and causes it to be communicated by the tags on the networks to the gateway tag. The gateway tag in turn communicates the ID to the computer 14 which retrieves the appropriate information. The information is routed back to the PDA through the gateway and the asset tags.

If a computer is unavailable, the user determines the assets contents from its package labeling, and enters the information directly into the PDA.

Once activated, the tag begins performing the functions described above. The location function of the tag interrogates the helper tag on the nearest conveyance for its location. The helper tag retrieves the information from the attached asset tag, and informs the tag of the location.

The tag then measures its distance from the equipment tag, in the manner described above. The tag repeats this measurement several times, and checks whether the distance has changed beyond a set error tolerance. If it has, the tag determines that it is not moving with the conveyance. If the distance has not changed, the tag continues to interrogate the equipment tag until it has enough measurements to accurately measure its distance from the conveyance.

The tag repeats this process for every conveyance in its vicinity, and uses the measured ranges and the location of the conveyances to determine the asset's location.

- 120 -

Once the tag knows its own location, it can determine the location of other tags, monitoring devices and gateways in the warehouse by determining its relative distance from these items.

5 The network formation unit in the tag periodically interrogates other asset tags, monitoring devices, gateways and GPS sensors. The tag measures its distance from the other tags and forms an optimal network with the tags, as described above. The tags also define
10 relay points between asset tags, gateways and monitoring devices, and define sentinels for communication with other networks.

Whenever a new tag is added to the network or removed from it, the gateway is informed. The gateway in
15 turn notifies the inventory computer of the change. In this manner, the inventory computer is notified of all movements of the inventory. If certain movements are expected, the computer notifies the gateway which passes the information onto the asset tags. Once notified, the
20 asset tags disregard the planned movements, and the inventory computer is notified only of unplanned movements.

The tags also perform all the self-assessment subfunctions. For example, when a tag recognizes that
25 its battery is failing, it notifies the gateway or a monitoring device and powers down. If the tag is a sentinel or relay point, a tag designated as a netmaster chooses a replacement for the tag from the nearby tags. Warehouse personnel receive set and pick lists from the
30 gateway tags on their monitoring devices. The personnel uses the PDA in the device to interrogate the network and locate the items on the list. The PDA may rearrange the list to minimize the path traveled by the personnel. The computer also notifies each tag within the network of the
35 set and pick list. After each set or pick the tags in

- 121 -

the network determine whether the item has been properly placed, and notifies the personnel through their PDAs of any error. These functions are omitted if the places are at the discretion of the personnel.

5 (c) Other embodiments

In other embodiments, tags are set up at the entrance to the warehouse or storage lot. Tags are affixed to the item (for example, motor vehicles, in particular cases) at the factory where they are produced. 10 The tag at the entrance of the storage lot monitors the vehicles and records their departure from the lot. Once on a dealer's lot, the cars form a network. The tag at the entrance can be replaced with a gateway and long distance communication capability provided on the car 15 itself or on a truck transporting the cars to and from the lot.

In some embodiments, the self-assessment functions of the asset tags are omitted. A client computer with inventory and communication software may be 20 connected to the inventory computer.

C. Warehouses with Hazardous Materials

Referring to Fig. 4, an application 10₂ for monitoring a warehouse storing hazardous materials is identical to the application 10₁ described above, but with 25 the addition of an interrogate and respond module 22 on an outside wall of the warehouse.

The module is connected to a serial port of inventory computer 14, and is powered by a battery with a built-in recharger connected to an external power source 30 at 110 VAC. The module may thus function during a power outage.

Warehouse personnel periodically instruct the inventory computer to download the warehouse contents in the module. Alternatively, the computer automatically 35 downloads the information without instruction from an

- 122 -

operator. When a fireman interrogates the module with a police radar, the module transmits the downloaded data in synthesized voice, followed by a data stream, over the fireman's communication frequency.

5 Alternatively, the module is in communication with a gateway tag in the warehouse. The gateway tag in turn communicates with the tags mounted on the items. When interrogated, the module causes the gateway to communicate with the asset tags and learn their contents.
10 The gateway then relays the information to the module.

The tags inform the gateway of their contents and location in increasing order of their ID. If the tags are fixed to containers, they communicate with tags on the assets inside the containers to ascertain their
15 contents.

In another embodiment, the module is connected to a dedicated response computer in the warehouse. The computer is equipped with a gateway tag in communication with the asset tags.

20 The dedicated computer periodically queries the assets through the gateway tag for the asset contents, and verifies the information. The computer arranges the information in an order for processing, and summarizes the data. For example, if 50 drums of a particular
25 material are located in one area of the warehouse, the computer replaces information relating the contents of each drum with a statement indicating that the fifty drums contain the same material. A higher level summary identifies only the quantity of materials in the
30 warehouse corresponding to each hazard class.

The computer responds to an interrogation from the module with information that is divided into sets. For example, a summary of the warehouse contents forms one set of data. More detailed information concerning a
35 particular area of the warehouse forms a subsequent set.

- 123 -

The module transmits one set of information to the fireman in synthesized voice, and follows it with a corresponding data stream. The module repeats this process for each set of data. The fireman may thus
5 display each set of data on a monitor to avoid a misunderstanding.

The dedicated computer is appropriate for large warehouses with a complex system of assets, since the dedicated computer organizes a large amount of
10 information before sending it to the module.

In another embodiment, the warehouse does not include any of the features of application 10₁, and data concerning the contents of the warehouse is entered manually on the computer and communicated to the module.
15 In all the embodiments described above, the monitoring device and the tags on the conveyances (location tags) may be omitted.

D. Railroad shipping

(a) System Design

Referring to Figs. 5 and 66, a railroad
20 application 10, monitors both the cargo on a train and the train's rolling stock (cars and locomotives). A locomotive 12₆ on the train is equipped with a monitoring device 15 comprising a PDA connected to a communications
25 controller with a boot and holster, as described above. The device is also equipped with a panic button, described above. A GPS sensor is mounted on the locomotive for communication with the controller. The locomotive also has an asset tag 16₆' connected to
30 operations center 13 through a cellular link 3005.

Railroad personnel in a loading or unloading area are also provided with an identical monitoring device. Each PDA is recharged with a standard charging cradle.

- 124 -

An asset tag connected to a helper tag (together referred to as a "rolling stock tag") is placed on cars and trailers in the train. The rolling stock tags have a larger energy source than those used in the applications described above, and have a communication range spanning more than one element of rolling stock. It is therefore not necessary to supply each car with a tag. Each piece of rolling stock is also equipped with one or more of the following sensors.

10 Wheel bearing sensors are placed close enough to a wheel bearing on the rolling stock to measure its temperature and communicate an exceptional temperature to the locomotive tag via the rolling stock tag. Brake release sensors monitor the rolling-stock brakes and
15 communicate a non-released brake condition to the locomotive tag via the rolling stock tag. Door sensors are mated with an asset tag and are in communication with the rolling stock tag. A reefer sensor is mated with an asset tag and monitors the fuel level, refrigeration
20 operation and reefer temperature for the rolling stock.

 Items 12, in the cars have asset tags connected to helper tags which communicate with the monitoring device 15. The helper tag and asset tag perform all the functions described above.

25 Selected rolling stock is equipped with a module 22 to allow an operator to interrogate the stock regarding its contents from a distance. The module is mounted outside a car and is coupled to a rolling stock tag located in the car. An operator may deactivate the
30 module through a PDA in the monitoring device.

(b) System operation

The system provides a means for efficiently searching for misrouted containers, and for providing an easily accessed compilation of the contents of the train.

- 125 -

During loading, an operator uses the PDA in the monitoring device to communicate with the rolling stock tag. The operator indicates to the tag which containers are being loaded onto the car. Similarly, the operator uses the PDA to relay the identity of the car to the tags on each container.

In this manner, the rolling stock tag and the asset tags are informed of one another's identities, and may begin communicating. The rolling stock tag selects a communications channel for communication with the asset tags and provides the identity of the channel to the asset tags. If excessive interference occurs on the channel, the rolling stock tag selects a new channel and informs the asset tags of the change.

After the train is loaded, the conductor notifies the operations center through the PDA of his identity. The center validates the operator's password and informs him of route information. Clients 3006 of the railroad 3008 (shippers and recipients of the cargo) interact either with the railroad directly or with the operations center.

The asset tags, tags on each piece of rolling stock and the sensors each perform the network formation subfunctions described above. Each rolling stock tag selects a communication channel for the sensors on the rolling stock, and communicate that choice to each sensor. If the tags subsequently detect excessive interference on the channel, they select another unused channel and instruct the sensors of the change. The sensors respond to instructions regarding the channel only from the tag for the rolling stock on which the sensors are located. The wheel bearing sensors periodically measure the temperature of the wheel bearings and communicates the temperature to a local asset tag. The tag compares the temperature to a

- 126 -

predetermined threshold. If the threshold is exceeded, the tag forwards a message to the nearest rolling stock tag. The rolling stock tag continues to communicate this message along the rolling stock tags until it reaches the locomotive tag. The locomotive tag forwards the message to the monitoring device, which, at the option of the railroad, sends the message to the railroad through the operations center.

The brake release sensors are activated by a local asset tag when the tag receives a brake release message through the network from the locomotive tag. After a preset delay time, the sensor determines whether the brake was actually released. The delay time allows sufficient time for the compressed air and mechanical braking system to operate. If the brakes have not been released, the tag notifies the monitoring device through the network that the brake is malfunctioning.

The door sensors receive a message through the network or through a monitoring device carried by an operator loading the train. The message may indicate either that the door need not be monitored, or that the door is open or closed and should remain that way. If the door violates these instructions, the tag sends a message through the network to the monitoring device in the locomotive. At the railroad's option, the monitoring device forwards this message to the railroad. The reefer sensors notify the monitoring device whenever a sensed value exceeds a preset threshold. The device may alert the railroad of this condition.

Each sensor also includes a test cycle, redundant sensors, or a calibration cycle that activates the sensor to determine whether it is still operational. The asset tags periodically monitor the sensors to determine whether they are still present and active.

- 127 -

The monitoring device on the locomotive or another device may deactivate any sensor to preserve energy. For example, the monitoring device may deactivate all the brake sensors once the train is in
5 motion or is stopped. The tag checks the ID of each device sending a deactivate message so that it is deactivated only by a valid device. The device can subsequently reactivate any sensor.

When a rolling stock tag is isolated from
10 other rolling stock tags, for example when the rolling stock is being unloaded or loaded, the rolling stock tag no longer functions as part of the network. Instead, an operator communicates with the tag using a PDA in a local monitoring device. If the stock is to be monitored while
15 unattended, a gateway tag is placed in the vicinity of the rolling stock. The tag networks with the gateway tag to notify a nearby computer of any changes.

The tags perform additional functions described in the warehousing application above. For
20 example, the tags notify the monitoring device in the locomotive through the network whenever a tag disappears. The device monitors the location of the train relative to an expected route, and notifies the center of any unacceptable deviations. Similarly the center may
25 communicate a new route to the device in the event of weather delays, etc.

The railroad, shipper or recipient may contact the center to locate a particular asset, in a manner analogous to that described above.

30 In the event of an emergency, the conductor activates and confirms the panic button. The monitoring device informs the center of the emergency, which takes action in a manner analogous to that described above. Similarly, if the operator is incapacitated, a policeman

- 128 -

can interrogate the module outside a car, or interrogate an asset directly.

(c) Other embodiments

In some embodiments, the operations center is
5 omitted, and the railroad communicates with the train's monitoring device directly. In this scenario, all communications take place through a distributed system of monitoring devices.

If the center is included, the client
10 communication, inventory, scheduling and accounting software described above may be omitted. In addition, the Internet link, UPS, generator, and redundant backups may be removed from the center.

In other embodiments, a client is provided
15 with a computer for communicating with the operations center. The computer has inventory software for processing inventory information relayed by the center.

In some embodiments, the sensors in the
20 locomotive and the holster and boot in the communications controller are omitted.

In some embodiments, the asset tags are
excluded, and the system only monitors the tags indirectly through the door sensor. If the tags are included, they may function without a helper tag. In
25 addition, the rolling stock tags are optional if the asset tags may directly communicate with asset tags in other cars.

In other embodiments, the attachment
subfunction and the active subfunction in the self-
30 assessment function are omitted from the asset tags.

E. Intermodal shipping through a port

Referring to Fig. 7, containers typically arrive
at a port by train 12₆' or truck 12₆. The containers are unloaded, placed on a flat truck, driven to a holding
35 area 64 and removed from the truck. A flat truck then

- 129 -

moves the containers to a ship 62 when it is ready for loading. The process occurs in reverse when a ship arrives at the port. A computer 14 in the management offices of the port monitor the containers and is in
5 communication with a remote operations center (Fig. 1).

Each container is equipped with an asset tag, and the trains and trucks are equipped as described in sections XIIA and XIID above. Each ship and holding area is treated as a warehouse and equipped with pseudolites,
10 gateways and inventory computers, as described above in Section XIIB.

If the size of the holding area or ship exceeds that of a typical warehouse, the communication range of the tags is extended by placing additional relay tags in
15 strategic locations around the ship and holding area. Each relay tag includes an asset tag connected to a helper tag similar to the pair mounted on the trailer in the hazardous waste application, but without the interrogate and respond capability. Each helper tag is
20 powered by a 110 VAC external power source.

Alternatively, the ship or holding area may be divided into areas and each area supplied with a gateway. The gateway is connected to the port management offices by a hard wire or long range wireless connection.
25 Combinations of the relay tags and gateways may also be used.

If the management offices wish to monitor the containers on a ship at sea, the gateways on the ship are connected by a global communications device, such as
30 InMarSat, in lieu of a cellular telephone connection.

Other embodiments of this application include the variations of the warehousing and transportation applications described above.

F. Bonded carriers

35 (a) System design

- 130 -

Bonded carriers operate conveyances that move containers and semi-trailers between national borders or storage areas near a border. Customs officials typically require that the conveyances prevent access to their contents. Additional security is provided by requiring the conveyances to follow specific routes over a specified time frame.

Referring to Fig. 6, each container 12₁ in the conveyance is equipped with an asset tag 16₁ connected to a seal 20', described in connection with Fig. 15. The asset tag and seal include all the units described above except the location and network formation units.

Each container is equipped with a second asset tag and seal. The second seal does not carry a GPS sensor, and has a smaller battery.

A remote operations center communicates with the carriers and the tags. An operator is provided with a monitoring device for communicating with the tags and center. A standard charging cradle is provided for the PDA in the monitoring device. Communications between the center and the tags are encrypted with a public/private key encryption system.

(b) System operation

As described in Section V, an operator attaches each seal to a container by threading an aircraft cable through a latch on the container, and placing the ends of the cable in two quick connect ports of the seal. An optical fiber is similarly threaded through the latch, and its ends inserted in a tamper detector in the seal.

In use, the carrier contacts the operations center, identifies a conveyance and its tags and describes a scheduled route for the conveyance. The carrier also indicates the carrier personnel who should be notified in case a deviation from the route occurs.

- 131 -

The center contacts the seals over the cellular phone and downloads the route information provided by the carrier.

The asset tags attached to the seals on each
5 container periodically power themselves on and attempt to communicate with one another. If a communication is successfully established, the asset tag attached to the first seal activates the GPS sensor and queries the sensor for its position and velocity, and time of day.
10 The first seal relays this information to the second seal on the container.

The asset tag attached to each seal then compares its location to an expected location derived from the route information. The seals report any
15 deviation to the center, and the attached asset tags reset their clocks and power down. The operations center periodically recalculates the expected route and downloads the route to the seals. For example, if a delay is expected due to snow, the center modifies the
20 route accordingly.

The center randomly contacts the seals over the cellular phone when the attached asset tags are powered down. When the asset tags wake up, they obtain location information from the GPS sensor and cause the
25 seals to relay the information to the center. The information includes the position of the sensor, its velocity and the absolute time.

The center recalculates the expected route from the information obtained from the seals. This
30 corrects for small modifications of the route that need not be communicated to the center by the seals.

The center also examines the route and calculates a likelihood that the conveyance will deviate from its expected route. The center then optimizes the
35 power on period for the seals based on the likelihood.

- 132 -

Tag energy is further conserved by deactivating the asset tags between trips. If the seals are unable to establish communication for example, because of equipment failure or tampering, a seal notifies the operations center. If the asset tag attached to the first seal is functioning, it activates the GPS sensor to determine its position and velocity and the time of day, and also reports this information to the center. If only the second seal is active, it notifies the center of its last known position, velocity and time. The seals continually transmit an IR beam through the fiber optic cable. If a seal is tampered with, that is, if the fiber is cut or damaged, the IR receiver in the tamper detector will not receive the beam through the fiber. The receiver immediately notifies the second seal on the container and calls the operations center. If the center does not confirm the call within a prespecified time, the seal notifies the second seal which attempts to make the communication. In applications where the containers include valuable cargo, a seal and asset tag are attached to each container. A module 22 is also mounted on the outside of the conveyance and is in communication with an interrogator inside the conveyance.

When the module is activated, it verifies the existence of each tag attached to a container. If the module is unable to establish communication with a tag, it notifies the interrogator which alerts the operations center.

In other embodiments, a door sensor is mounted near the door of the conveyance. When the sensor determines that the door is improperly opened, it alerts a seal on a container. This provides an indirect monitoring of the contents of the conveyance. If the cargo is refrigerated, a temperature sensor is placed near the containers. If the sensor notes that the

- 133 -

temperature is outside an acceptable range, it notifies a tag.

(c) Other embodiments

In some embodiments, the seals and asset tags
5 include the location unit and network formation unit,
discussed in connection with Fig. 23.

A related application includes avoiding or
reducing theft in shipping and storing pipes. A tag with
a tamper detectors is placed on one end of a pipe and
10 connected by a repeater, for example, an asset tag
relaying IR communications between the tags at the other
end of the pipe. One end of an optical fiber is placed
in one tag and run through the pipe. An opposite end of
the fiber is connected to the second tag. The first tag
15 transmits a signal through the fiber to the second tag.
The second tag relays the signal as a repeater, which
sends the signal to the first tag. When either tag
detects an absence of a signal at its receiver, it
notifies a gateway tag. G. Clandestine Tracking

20 Referring to Fig. 9, application 10, attaches a
tag 16' to an asset such that the tag is difficult to
detect. This is accomplished by for example, mounting
the tag behind the package labelling for the asset.
Smaller tags are particularly useful in this application;
25 as described above, the tag size is reduced to 2 cm in
diameter and 1/4 cm in thickness when a lower life
battery is employed.

Without a monitoring device, a smuggler of the
asset will be unable to electronically detect the
30 presence of the tag. By contrast, once the asset is
loaded onto a conveyance with a gateway tag, the gateway
will detect the tag and report the presence of the asset
to a monitoring device or operations center. Similarly,
personnel will be able to activate the gateway when in
35 the gateway's range to detect the presence of the asset.

- 134 -

If the asset is being carried in the open, or if a door to the conveyance is opened, the tag can be detected by tags on other conveyances in the area, or by the monitoring devices of nearby personnel. A charging
5 cradle is provided for these monitoring devices. The tags perform the seek others subfunction, but other network formation and location functions are optional. Self-assessment functions are also optional, but the read and write subfunction is required.

10 This application is particularly useful for tracking controlled substances and other illegal materials.

Other embodiments are within the following claims.

Appendix I: Calculation of f in tetrahedron

Referring to Fig. 67, in the planar tetrahedron defined by tags A,B,C,D, the distances between tags A, B and C are known. D₁ is a location for the fourth tag D that minimizes the range between C and D (CD₁ or f_{min}), and D₂ is a location for tag D that maximizes the distance between C and D (CD₂ or f_{max}).

The equations that follow define f_{min} and f_{max} in terms of the ranges between tags A, B and C. Table I defines the notation used in the equations.

Table I

\overline{AB}	The length of segment AB.
$\overline{\overline{AB}}$	The length of segment AB, with a sign factor (i.e. it can be negative)
$\angle BAC$	The angle defined by segments AB and AC at point A.
$\triangle BAC$	The triangle defined A, B and C.
$\angle \triangle ABC, \triangle ABD$	The angle between triangles ABC and ABD.

$$\overline{\overline{AR_2}} = \frac{\overline{BD^2} - \overline{AB^2} - \overline{AD^2}}{2 \overline{AB}}$$

$$\overline{DR_2} = \sqrt{\overline{AD^2} - \overline{AR_2^2}}$$

$$\overline{\overline{AR_1}} = \frac{\overline{BC^2} - \overline{AB^2} - \overline{AC^2}}{2 \overline{AB}}$$

$$\overline{CR_1} = \sqrt{\overline{AC^2} - \overline{AR_1^2}}$$

$$x = \overline{\overline{AR_1}} - \overline{\overline{AR_2}}$$

$$y_{min} = \overline{DR_2} - \overline{CR_1}$$

$$y_{max} = \overline{DR_2} + \overline{CR_1}$$

$$\overline{CD_1} = \sqrt{x^2 + y_{min}^2}$$

$$\overline{CD_2} = \sqrt{x^2 + y_{max}^2}$$

$$\overline{CD_1} \leq \overline{CD} \leq \overline{CD_2}$$

Appendix II:

Calculation of j_1 and j_2 in five tag group

Referring to Fig. 68, tags A, B, C and D define a tetrahedron. Tag E (not shown) defines a second tetrahedron that shares triangle ABC with the first tetrahedron.

The altitude DQ_1 is perpendicular to triangle ABC, and point Q_1 lies on the triangle ABC. Q_1R_1 is perpendicular to AB. DR_1 is the altitude of triangle ABD and CS is the altitude of triangle ABC. Points Q_2 and R_2 are placed analogously to Q_1 and R_1 in the second tetrahedron.

The following equations define the relationships between each side of the triangle. The final equation defines two solutions for the range DE (j_1 and j_2). This equation can be manipulated to derive an expression for each remaining side in terms of the nine other sides.

Table I defines the notation used in the equations.

$DR_1 \perp AB$
 $CS \perp AB$
 $Q_1R_1 \perp AB$
 $DQ \perp \Delta ABC$

$$\cos(\angle BAD) = \frac{BD^2 - AB^2 - AD^2}{2 AD AB}$$

$$\cos(\angle BAC) = \frac{BC^2 - AB^2 - AC^2}{2 AC AB}$$

$$\overline{AS} = \overline{AC} \cos(\angle BAC)$$

$$\overline{CS} = \overline{AC} \sin(\angle BAC) = + \sqrt{AC^2 - \overline{AS}^2}$$

$$\overline{AR_1} = \overline{AD} \cos(\angle BAD)$$

$$\overline{DR_1} = \overline{AD} \sin(\angle BAD) = + \sqrt{AD^2 - \overline{AR_1}^2}$$

$$\overline{R_1S} = |\overline{AS} - \overline{AR_1}|$$

$$\alpha_1 = \angle \Delta ABC, \Delta ABD$$

$$\overline{CD}^2 = \overline{R_1S}^2 + (\overline{DR_1}^2 + \overline{CS}^2 - 2\overline{DR_1} \overline{CS} \cos \alpha_1)$$

$$\cos \alpha_1 = \frac{\overline{R_1S}^2 + \overline{DR_1}^2 + \overline{CS}^2 - \overline{CD}^2}{2 \overline{DR_1} \overline{CS}}$$

$$\overline{Q_1R_1} = \overline{DR_1} \cos \alpha_1$$

$$\overline{DQ_1} = \overline{DR_1} \sin \alpha_1 = \pm \sqrt{\overline{DR_1}^2 - \overline{Q_1R_1}^2}$$

$$x_1 = \overline{AR_1}$$

$$y_1 = \overline{Q_1R_1}$$

$$z_1 = \overline{DQ_1}$$

$ER_2 \perp AB$
 $Q_2R_2 \perp AB$
 $EQ \perp \Delta ABC$

$$\cos(\angle BAE) = \frac{BE^2 - AB^2 - AE^2}{2 AE AB}$$

$$\overline{AS} = \overline{AC} \cos(\angle BAC)$$

$$\overline{CS} = \overline{AC} \sin(\angle BAC) = \sqrt{AC^2 - \overline{AS}^2}$$

$$\overline{AR_2} = \overline{AE} \cos(\angle BAE)$$

$$\overline{ER_2} = \overline{AE} \sin(\angle BAE) = \sqrt{AE^2 - \overline{AR_2}^2}$$

$$\overline{R_2S} = |\overline{AS} - \overline{AR_2}|$$

$$\alpha_2 = \angle \Delta ABC, \Delta ABE$$

$$\overline{CE}^2 = \overline{R_2S}^2 + (\overline{ER_2}^2 + \overline{CS}^2 - 2\overline{ER_2} \overline{CS} \cos \alpha_2)$$

$$\cos \alpha_2 = \frac{\overline{R_2S}^2 + \overline{ER_2}^2 + \overline{CS}^2 - \overline{CE}^2}{2 \overline{ER_2} \overline{CS}}$$

$$\overline{Q_2R_2} = \overline{ER_2} \cos \alpha_2$$

$$\overline{EQ_2} = \overline{ER_2} \sin \alpha_2 = \pm \sqrt{\overline{ER_2}^2 - \overline{Q_2R_2}^2}$$

$$x_2 = \overline{AR_2}$$

$$y_2 = \overline{Q_2R_2}$$

$$z_2 = \overline{EQ_2}$$

$$\overline{DE} = + \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (|z_1| \pm |z_2|)^2}$$

- 137 -

Appendix III: Deriving the covariance matrix

For a given range (m), a particular grouping of five coplanar tags containing the range is chosen. The five tags define two tetrahedrons sharing a common side, as shown in Fig 68.

An estimate m_1 for the range is derived by solving for the range from the nine other ranges (a,b,c,d,e,f,g,h,i) in that grouping, such that

$$m_1 = f(a,b,c,d,e,f,g,h,i)$$

The function $f(a,b,c,d,e,f,g,h,i)$ is derived from the Equations given in Appendix II, and depends on the location of the range m within the pair of tetrahedrons.

A similar equation provides other estimates m_2, m_3, \dots using other grouping of five tags containing the range m.

As described above, the equation yields two values for m_1 (for example, j_1 and j_2 for side DE). The processor chooses the value closest to the range originally measured for m. Alternatively, when the range m is invalid, a value is chosen from each pair of values m_1, m_2, \dots , such that the chosen values have a minimum variance.

The derivation of the variances of each estimate (and their covariances) is demonstrated from a simplified case, in which

$$\begin{aligned} m_1 &= a \\ m_2 &= f(b,c,d) \\ m_3 &= f(c,d,e) \end{aligned}$$

The function $f(b,c,d)$ and $f(c,d,e)$ are modified to return only the desired value of the estimate. The general case, where each estimate is a function of nine ranges, and where there are a total of N estimates, can be easily derived by generalizing the following discussion.

The ranges (a,b,c,d,e) are statistically independent, and are associated with a Gaussian error (delta) with variance (sigma). Each estimate m_1, m_2, m_3 is linearized to create x_1, x_2, x_3 such that

$$\begin{aligned} x_1 &= a + \Delta a \\ x_2 &= f(b, c, d) + \frac{\partial f(b, c, d)}{\partial b} \Delta b + \frac{\partial f(b, c, d)}{\partial d} \Delta d + \frac{\partial f(b, c, d)}{\partial c} \Delta c \quad (34) \\ x_3 &= f(c, d, e) + \frac{\partial f(c, d, e)}{\partial c} \Delta c + \frac{\partial f(c, d, e)}{\partial d} \Delta d + \frac{\partial f(c, d, e)}{\partial e} \Delta e \end{aligned}$$

5 The variances of each linearized estimate is given by:

$$\begin{aligned} \sigma_{11}^2 &= \sigma_a^2 \\ \sigma_{22}^2 &= \left(\frac{\partial f(b, c, d)}{\partial b} \right)^2 \sigma_b^2 + \left(\frac{\partial f(b, c, d)}{\partial d} \right)^2 \sigma_d^2 + \left(\frac{\partial f(b, c, d)}{\partial c} \right)^2 \sigma_c^2 \quad (35) \\ \sigma_{33}^2 &= \left(\frac{\partial f(c, d, e)}{\partial c} \right)^2 \sigma_c^2 + \left(\frac{\partial f(c, d, e)}{\partial d} \right)^2 \sigma_d^2 + \left(\frac{\partial f(c, d, e)}{\partial e} \right)^2 \sigma_e^2 \end{aligned}$$

Because m_1 is only dependent on a and m_2 and m_3 are independent of a, the covariances between them are zero. The covariance between x_2 and x_3 is given by the expectation:

$$\sigma_{23}^2 = E[(x_2 - m_2)(x_3 - m_3)] \quad (36)$$

10 Equations 36 and 34 together yield:

$$\sigma_{23}^2 = \frac{\partial f(b, c, d)}{\partial c} \frac{\partial f(c, d, e)}{\partial c} \sigma_c^2 + \frac{\partial f(b, c, d)}{\partial d} \frac{\partial f(c, d, e)}{\partial d} \sigma_d^2 \quad (37)$$

The covariance matrix is thus given by:

$$\Sigma = \begin{vmatrix} \sigma_{11}^2 & 0 & 0 \\ 0 & \sigma_{22}^2 & \sigma_{23}^2 \\ 0 & \sigma_{23}^2 & \sigma_{33}^2 \end{vmatrix} \quad (38)$$

The inverse Q of the covariance matrix is defined by the matrix (q_{ij}) .

For Q to exist, it is essential that the
 15 covariance matrix be nonsingular. This is guaranteed by selecting groups of five tags (5-tuples) that contain the

- 139 -

range (m), such that each selected group includes a range absent from the previously selected groups.

Referring to Fig. 69, the subfunction selects the 5-tuples by first defining the range (m) being refined
5 (step 4000). The subfunction selects an arbitrary set of three tags (a 3-tuple) that does not include either tag associated with the range m (step 4002), and combines the 3-tuple with the tags defining the range m to form a 5-tuple (step 4004).

10 The subfunction next finds a 3-tuple that contains at least one range not included in any previously selected 3-tuple (step 4006). If no such 3-tuple exists (step 4008), the subfunction ends, and calculates Q for the set of 5-tuples selected up to that step. If no 5-
15 tuples have been selected, the subfunction does not refine the range m. The subfunction then returns to step 4000, and selects another range to be refined (step 4009). When all the ranges have been refined, the process ends (step 4010)

20 If an appropriate 3-tuple exists at step 4006, the subfunction combines the latest selected 3-tuple with the tags defining the range m (step 4011) to form a new 5-tuple. The subfunction then adds the latest 5-tuple constructed by the subfunction to a list of 5-tuples to
25 be used in calculating Q (step 4012) and returns to step 4006.

After a range m has been refined by the subfunction, it is not used in subsequent range refinements. This is because the equations given above
30 are derived by assuming that the ranges are statistically uncorrelated. Once the ranges are refined, they become correlated with other ranges, and employing them would therefore be inconsistent with this assumption. The refined ranges are used for the first time by the compute
35 locations subfunction 506 (Fig. 25).

- 140 -

What is claimed is:

1. A system for monitoring and affecting actions on behalf of a group of objects, comprising:

5 a first electronic tag attached to an object in the group, comprising:

circuitry for communicating with a second tag, the circuitry communicating at least information pertaining to an object in the group, and

10 a memory connected to the circuitry and capable of storing the information, and

a second electronic tag attached to an object in the group, comprising

circuitry for communicating the information with the first tag,

15 a memory connected to the circuitry in the second tag and capable of storing the information,

wherein the information comprises a distance between the first tag and the second tag.

2. The system of claim 1, wherein the first tag
20 comprises a time measuring device connected to a controller for measuring a time duration required for a signal to travel between the first tag and the second tag.

3. The system of claim 2, wherein
25 the time measuring device in the first tag comprises a clocked counter which has a first count corresponding to the time of transmission of the signal and which subsequently has a second count corresponding to the time of recognition that said signal has been
30 received; and

- 141 -

said device uses the product of the clock period and the difference between said counts to determine the time duration.

4. The system of claim 3, wherein the clock
5 counter starts in response to the transmission of the signal and stops in response to the reception of said signal.

5. The system of claim 3, wherein the memory in
the first tag includes a program configuring the
10 controller in the first tag to execute a function when the controller reads the memory, the function causing the controller to determine the distance based on the time duration.

6. The system of claim 5, wherein the second tag
15 further comprises a second counter connected to the controller in the second tag, the counter in the first tag further comprising a mechanism for adjusting a frequency of the first-mentioned counter, the mechanism being coupled to the controller in the first tag, the
20 program in the first tag causing the controller in the first tag to employ the mechanism to offset the frequency of the first-mentioned counter from a frequency of the counter in the second tag.

7. The system of claim 6 wherein the program in
25 the first tag configures the controller to adjust the frequency of the counter in the first tag based on the distance.

8. The system of claim 7 wherein the program in
the first tag causes the controller in the first tag to

- 142 -

perform in a specified and predetermined manner the functions of:

adjusting the frequency of the counter in the first tag based on the frequency of the counter in the
5 second tag, and

enabling the circuitry in the first tag to transmit a number of symbols at an intersymbol period related to the frequency of the counter in the first tag.

9. The system of claim 8 wherein the program in
10 the first tag further causes the controller in the first tag to perform the function of selecting the number of symbols and the frequency of the counter in the first tag such that a number of clock cycles of the first tag during a transmission and receipt of the symbols is
15 substantially equal to a number of clock cycles of the clock of the second tag within an offset (n), n and the number of symbols having no common factors other than one.

10. The system of claim 2, wherein the second tag
20 further comprises a delay circuit connected between the circuitry and the controller in the second tag, the delay circuit causing the second tag to delay returning the signal to the first tag by a given time period.

11. A method of measuring a distance between two
25 objects, said method comprising the steps of:

transmitting from one of said objects a plurality of symbols, said transmissions occurring in a first series in which consecutive transmissions are separated by a first time interval which is a multiple of a first
30 clock period;

receiving each of said symbols at the other of said objects and re-transmitting each of said symbols

- 143 -

back to said one of said objects, said re-transmissions occurring in a second series in which consecutive re-transmissions are separated by a second time interval which is different from the first time interval and which
5 is a multiple of a second clock period which is different from the first clock period;

measuring a round-trip time corresponding to detection of receipt from said other of said objects of each of said symbols, each said round-trip time being a
10 multiple of said first clock period; and

determining said distance between said objects using a velocity of said symbols, said clock periods, at least one of said time intervals, and said round-trip times.

15 12. The method of claim 11, wherein one clock period is approximately equal to the product of said other clock period and a fraction; the numerator of said fraction is equal to the sum of two terms;

20 the first term of said sum is the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval;

the second term of said sum is a whole number that shares no multiplicative factors, other than 1, with said
25 number of symbols transmitted; and

the denominator of said fraction is equal to the product of the number of symbols transmitted and said number of first clock periods corresponding to said first time interval.

30 13. The method of claim 11, wherein the measured round-trip times vary and have a variation range equal to one first clock period.

- 144 -

14. A method of measuring a distance between two objects, said method comprising the steps of:

transmitting from one of said objects a plurality of symbols, said transmissions occurring in a series in
5 which consecutive transmissions are separated by a time interval which is a multiple of a first clock period;
returning each of said symbols from the other of said objects back to said one of said objects;
measuring a round-trip time corresponding to
10 detection of receipt from said other of said objects of each of said symbols, each said round-trip time corresponding to a multiple of a second clock period; and
determining said distance between said objects using a velocity of said symbols, said clock periods, at
15 least one of said time intervals, and said round-trip times.

15. The method of claim 14, further comprising a step of detecting each transmission using the second clock period, wherein the round-trip times correspond to
20 time intervals between detections of receipt and transmission of the symbols.

16. The method of claim 14, wherein one of the clock periods is approximately equal to the product of the other clock period and a fraction;
25 the numerator of said fraction is equal to the sum of two terms;
the first term of said sum is the product of the number of symbols transmitted and the number of first clock periods corresponding to the first time interval;
30 the second term of said sum is a whole number that shares no multiplicative factors, other than 1, with said number of symbols transmitted; and

- 145 -

the denominator of said fraction is equal to the product of the number of symbols transmitted and said number of first clock periods corresponding to said first time interval.

5 17. The method of claim 14, wherein the measured round-trip times vary and have a variation range equal to one first clock period.

18. A method of measuring a distance between two objects, said method comprising the steps of:

10 transmitting from one of said objects a first symbol;

 transmitting from said one of said objects a second symbol, said transmission of said second symbol being separated from said transmission of said first
15 symbol by a time interval which is a multiple of a first clock period;

 receiving said first symbol at the other of said objects and re-transmitting said first symbol back to said one of said objects;

20 receiving said second symbol at said other of said objects and re-transmitting said second symbol back to said one of said objects, said re-transmission of said second symbol being separated from said re-transmission of said first symbol by a second time interval which is a
25 multiple of a second different clock period;

 measuring a first round-trip time corresponding to detection of receipt from said other of said objects of said first symbol, said first round-trip time being a multiple of said first clock period;

30 measuring a second round-trip time corresponding to detection of receipt from said other of said objects of said second symbol, said second round-trip time being a multiple of said first clock period;

- 146 -

if said round-trip times are not different by at least one first clock period, repeating said transmitting, said receiving, said re-transmitting, and said round-trip measuring of said second symbol until
5 said round-trip times are different by one first clock period; and

determining the distance between said objects using a velocity of the symbols, the clock periods, the time intervals, and the round-trip times.

10 19. The method of claim 18, wherein the first and second round-trip times are original round-trip times, said original round-trip times being different by one first clock period, said method further comprising the steps of

15 (a) equating the clock periods;
(b) temporarily changing one of said clock periods for a period of time equal to a fraction of the time interval;

(c) determining a new round-trip time for a
20 symbol;

(d) comparing said new round-trip time to at least one of said original round-trip times; and

repeating (a) through (d) for a plurality of iterations, said fraction being different for each
25 iteration.

20. A method of measuring a distance between two objects, one of said objects being able to measure time with a resolution of a minimum period of time, said method comprising the steps of:

30 transmitting a plurality of symbols from said one of said objects;

receiving each of said symbols at the other of said objects and re-transmitting each of said symbols

- 147 -

back to said one of said objects after a delay period,
said delay period being different for each symbol and
being a multiple of a constant delay unit, said delay
unit being approximately equal to said minimum period
5 divided by the number of symbols transmitted;

measuring a round-trip time corresponding to
detection of receipt of each said symbol from said other
of said objects, each round-trip time corresponding to a
multiple of said minimum period; and

10 determining said distance between said objects
using a velocity of said symbols, said minimum period,
and at least one of said round-trip times.

21. A method of measuring a distance between two
objects, one of said objects being able to measure time
15 with a resolution of a minimum period of time, said
method comprising the steps of:

transmitting a plurality of symbols from said one
of said objects;

after a first delay period, detecting each
20 transmission, said first delay period being different for
each transmission and being a multiple of a constant
delay unit, said delay unit being approximately equal to
said minimum period divided by the number of symbols
transmitted;

25 returning each of said symbols at the other of
said objects back to said one of said objects;

after a second delay period, detecting the receipt
of each of said symbols from said other of said objects,
said second delay period being different for each receipt
30 and being a multiple of said delay unit;

measuring a time difference between detections of
transmission and receipt of each of said symbols, each
said time difference being a multiple of said minimum
period; and

- 148 -

determining said distance between said objects using a velocity of said symbols, said minimum period, and said time differences.

22. A method of measuring a distance between two
5 objects, one of said objects being able to measure time with a resolution of a minimum period of time, said method comprising the steps of:

transmitting a first symbol from said one of said
objects;

10 receiving said first symbol at the other of said objects and re-transmitting said first symbol back to said one of said objects;

measuring a first round-trip time corresponding to
detection of receipt from said other of said objects of
15 the first symbol, said first round-trip time being a multiple of said minimum period;

a) transmitting another symbol from said one of
said objects;

b) receiving said another symbol at said other of
20 said objects and re-transmitting said another symbol back to said one of said objects after a delay period, said delay period being a multiple of a delay unit, said delay unit being approximately equal to a fraction of said minimum period, the number of delay units in said delay
25 period being equal to the number of symbols previously retransmitted;

c) measuring another round-trip time corresponding
to detection of receipt from said other of said objects
of said another symbol, said another round-trip time
30 being a multiple of said minimum period;

if said first round-trip time and said another
round-trip time are equal, repeating a) through c) until
said round-trip times are different by one said minimum
period; and

- 149 -

determining said distance between said objects by using a velocity of said symbols, said minimum period, delay unit, and said round-trip times.

23. The method of claim 22, further comprising
5 reducing the delay unit to a smaller fraction of the minimum period and repeating a) through c).

24. Apparatus for measuring a distance between two objects, said apparatus comprising

a transmitter having a first clock for
10 transmitting from one of said objects a plurality of symbols, said transmissions occurring in a first series in which consecutive transmissions are separated by a first time interval which is a multiple of the period of said first clock;

15 a transponder having a second clock for receiving each of said symbols at the other of said objects and re-transmitting each of said symbols back to said one of said objects, said re-transmissions occurring in a second series in which consecutive re-transmissions are
20 separated by a second time interval which is different from said first time interval and which is a multiple of said second clock's period which is different from the period of said first clock;

a clock counter for measuring each round-trip time
25 corresponding to detection of receipt from said other of said objects of each of said symbols, each round-trip time being a multiple of said first clock period; and

a calculator for determining said distance between said objects using a velocity of said symbols, said clock
30 periods, at least one of said time intervals, and said round-trip times.

- 150 -

25. Apparatus for measuring a distance between two objects, said apparatus comprising:

a transmitter having a first clock for transmitting from one of said objects a plurality of
5 symbols, said transmissions occurring in a series in which consecutive transmissions are separated by a time interval which is a multiple of the period of said first clock;

a transponder for returning each of said symbols
10 from the other of said objects back to said one of said objects;

a counter having a second clock for measuring a round-trip time corresponding to detection of receipt from said other of said objects of each of said symbols,
15 each said round-trip time corresponding to a multiple of the period of said second clock; and

a calculator for determining said distance between said objects using a velocity of said symbols, said clock periods, said time interval, and said round-trip times.

20 26. Apparatus for measuring a distance between two objects, said apparatus comprising:

a transmitter for transmitting from one of said objects a first symbol and transmitting from said one of said objects a second symbol, said transmission of said
25 second symbol being separated from said transmission of said first symbol by a time interval which is a multiple of a first clock period;

a transponder for (a) receiving said first symbol at the other of said objects and re-transmitting said
30 first symbol back to said one of said objects and (b) receiving said second symbol at said other of said objects and re-transmitting said second symbol back to said one of said objects, said re-transmission of said second symbol being separated from said re-transmission

- 151 -

of said first symbol by a second time interval which is a multiple of a second different clock period;

a clock counter for (a) measuring a first round-trip time corresponding to detection of receipt from said other of said objects of said first symbol, said first round-trip time being a multiple of said first clock period, and (b) measuring a second round-trip time corresponding to detection of receipt from said other of said objects of said second symbol, said second round-trip time being a multiple of said first clock period;

a comparator for (a) testing if said round-trip times are not different by one first clock period, and (b) initiating repeating of said transmitting, said receiving, said re-transmitting, and said round-trip measuring of said second symbol until said round-trips are different by one first clock period; and

a calculator for determining the distance between said objects using a velocity of the symbols, the clock periods, the time intervals, and the round-trip times.

27. A method of measuring a distance between two objects, said method comprising the steps of:

transmitting from one of said objects a plurality of symbols, said transmissions occurring in a first series in which consecutive transmissions are separated by a first time interval which is a multiple of a first clock period, said symbols having a constant velocity;

receiving each of said symbols at the other of said objects and re-transmitting each of said symbols back to said one of said objects, said re-transmissions occurring in a second series in which consecutive re-transmissions are separated by a second time interval, said second time interval being different from the first time interval, said second time interval being a multiple of a second clock period which is different from the

- 152 -

first clock period, each said re-transmission thus occurring after a delay, said delay being different for each re-transmission and constituting the period of time between arrival of each said symbol at said other of said
5 objects and re-transmission of each said symbol, said delay having a nominal length which is a multiple of a delay unit, said delay unit being equal to the product of the number of first clock periods corresponding to the
10 first time interval and the difference between the clock periods, said delay unit also being equal to the second clock period divided by the number of symbols in the plurality of symbols, said delay having an actual length which is within one delay unit of said nominal length;

measuring a round-trip time corresponding to
15 detection of receipt from said other of said objects of each of said symbols, each said round-trip time being a multiple of said first clock period; and

determining said distance between said objects by multiplying said velocity of said symbols by half of the
20 difference between the average of all of said round-trip times and the first clock period.

28. The method of claim 27, wherein
the delay further comprises an extra delay which is a multiple of said second clock period; and
25 the distance is determined by multiplying said velocity by half of the difference between (a) the average of all of said round-trip times and (b) the sum of said extra delay and the first clock period.

29. A method of measuring a distance between two
30 objects, said method comprising the steps of:

transmitting from one of said objects a plurality of symbols, said transmissions occurring in a series in which consecutive transmissions are separated by a time

- 153 -

interval which is a multiple of a first clock period,
said symbols having a constant velocity;

detecting said transmissions using a second clock
period, said second clock period being different from
5 said first clock period, each said transmission detection
thus occurring after a delay, said delay being different
for each said detection and constituting the period of
time between each said transmission and said detection of
each said transmission, said delay having a nominal
10 length which is a multiple of a delay unit, said delay
unit being equal to the product of the number of first
clock periods corresponding to the time interval and the
difference between the clock periods, said delay unit
also being equal to the second clock period divided by
15 the number of symbols in the plurality of symbols, said
delay having an actual length which is within one delay
unit of said nominal length;

returning each of said symbols from the other of
said objects back to said one of said objects;

20 detecting the arrivals of said returning symbols
back at said one of said objects using said second clock
period, each said arrival detection thus occurring after
a second delay, said second delay being different for
each said detection and constituting the period of time
25 between each said arrival and said detection of each said
arrival, the nominal length of said second delay being a
multiple of said delay unit, said second delay having an
actual length which is within one delay unit of said
nominal length of said second delay;

30 measuring a round-trip time corresponding to
detection of receipt from said other of said objects of
each of said symbols, each said round-trip time
corresponding to a multiple of a second clock period; and

- 154 -

determining said distance between said objects by multiplying said velocity by half of the average of the round-trip times.

30. The method of claim 29, wherein
5 the symbols are returned from the other of said objects after an extra delay; and
the distance is determined by multiplying said velocity by half of the difference between the average of all of the round-trip times and said extra delay.

10 31. The method of claim 27, further comprising a step of offsetting the average of the round-trip times, wherein the offset is an amount of time equal to the second clock period divided by twice the product of (a) the number of symbols and (b) the number of first clock
15 periods corresponding to the first time interval.




FIG. 1




FIG. 2




FIG. 3




FIG. 5




FIG. 6





FIG. 7

8/59

FIG. 8






FIG. 9




FIG. 10

11/59




FIG. 11

12/59




FIG. 12




FIG. 14




FIG. 15




FIG. 16

17/59




FIG. 17

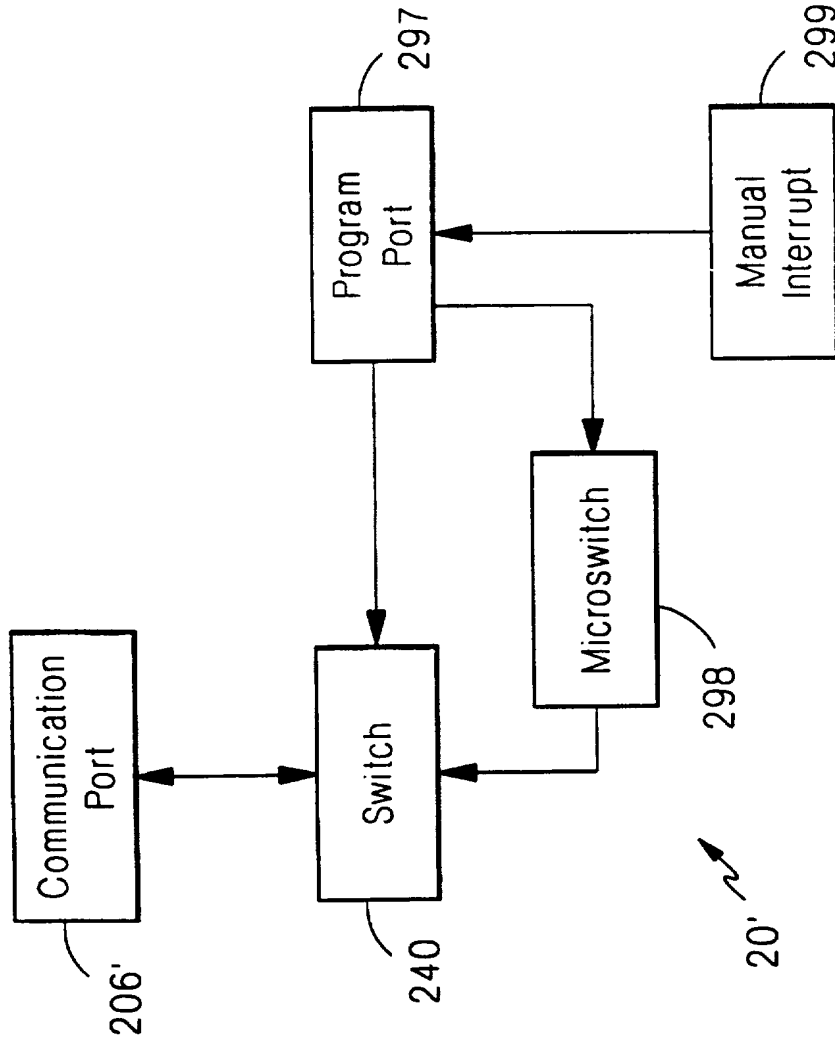


FIG. 19

20/59




FIG. 20
SUBSTITUTE SHEET (RULE 26)




FIG. 21

22/59

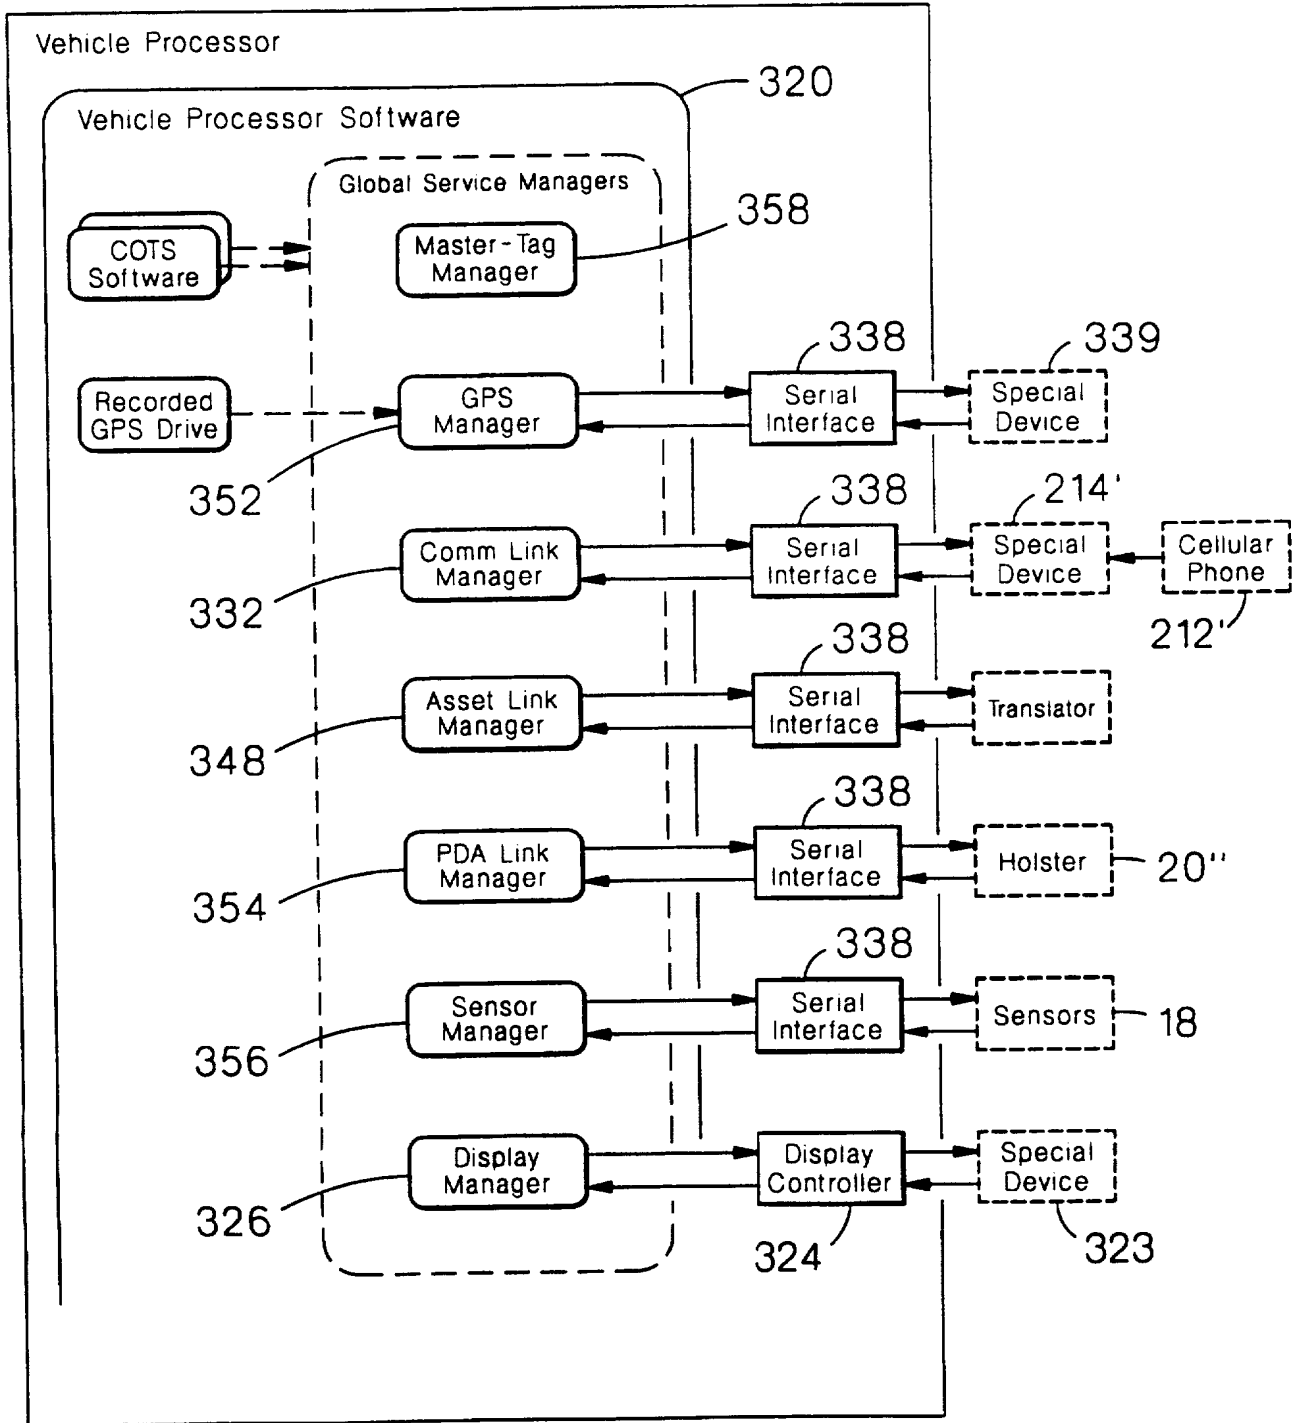


FIG. 22

23/59




FIG. 23

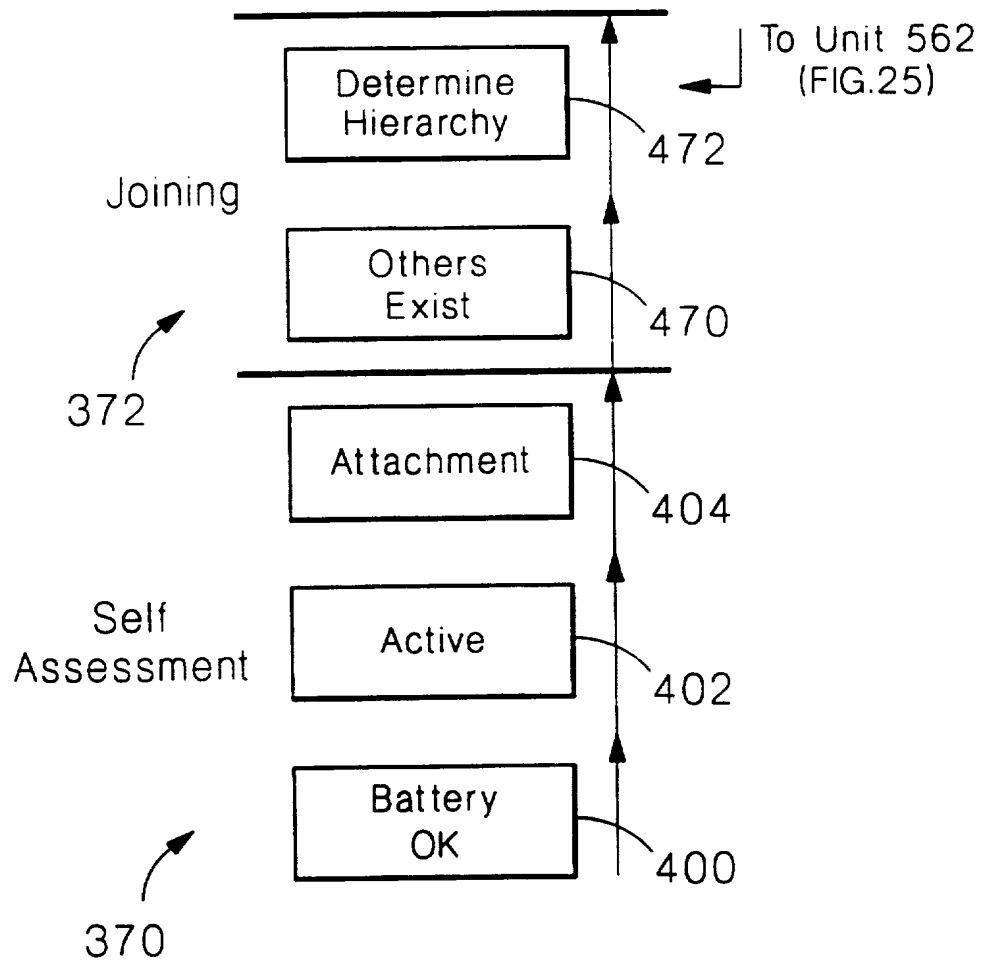


FIG. 24

25/59




FIG. 25

26/59

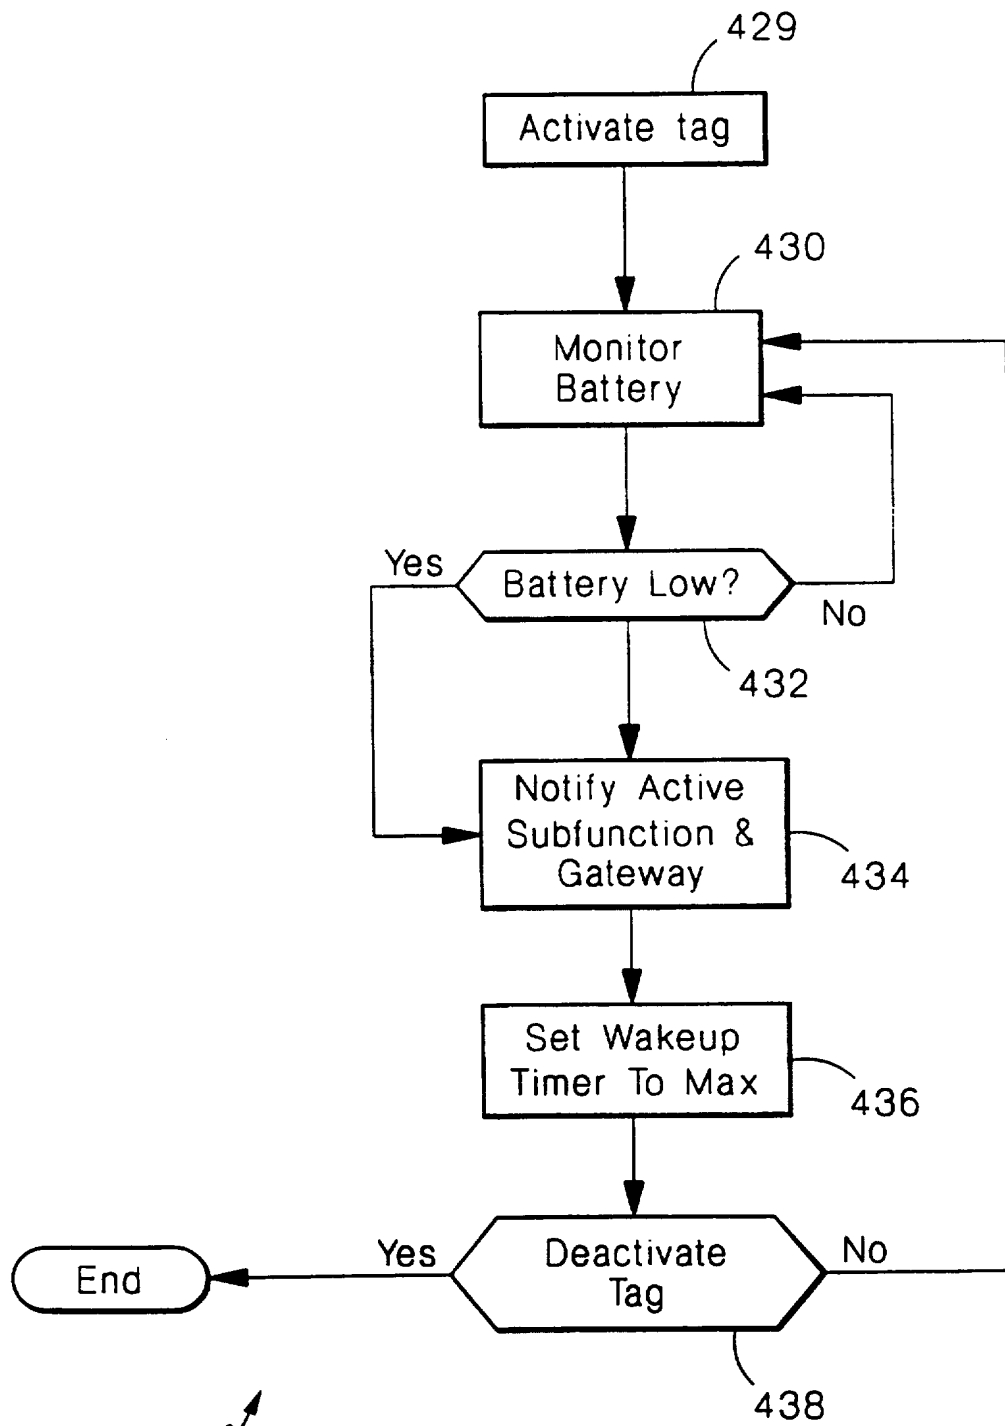


FIG. 26

27/59

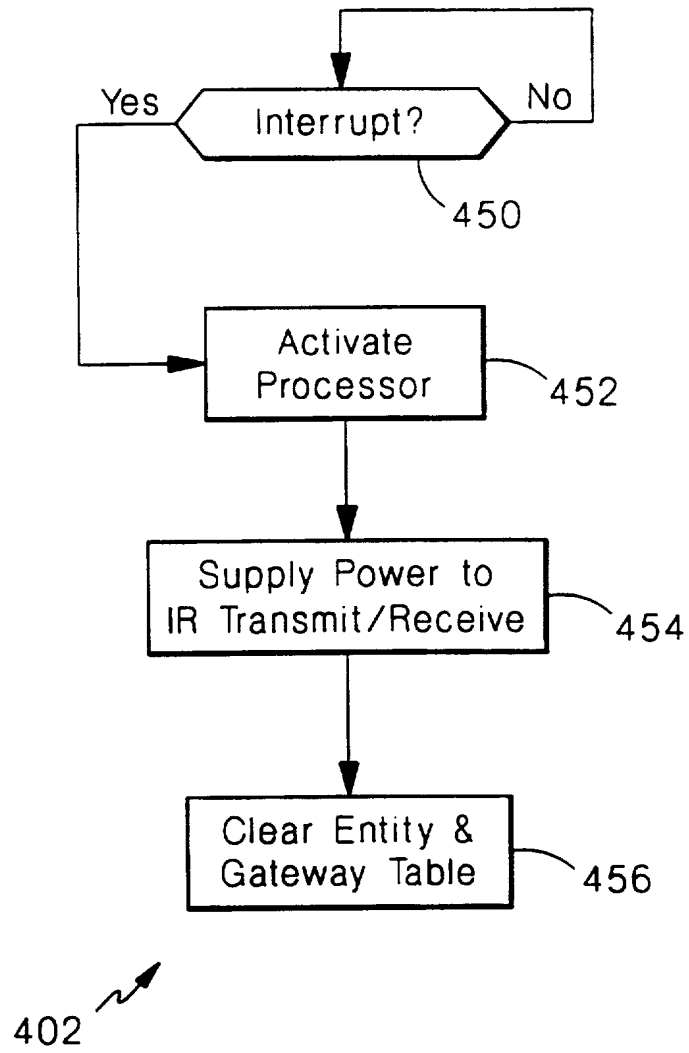
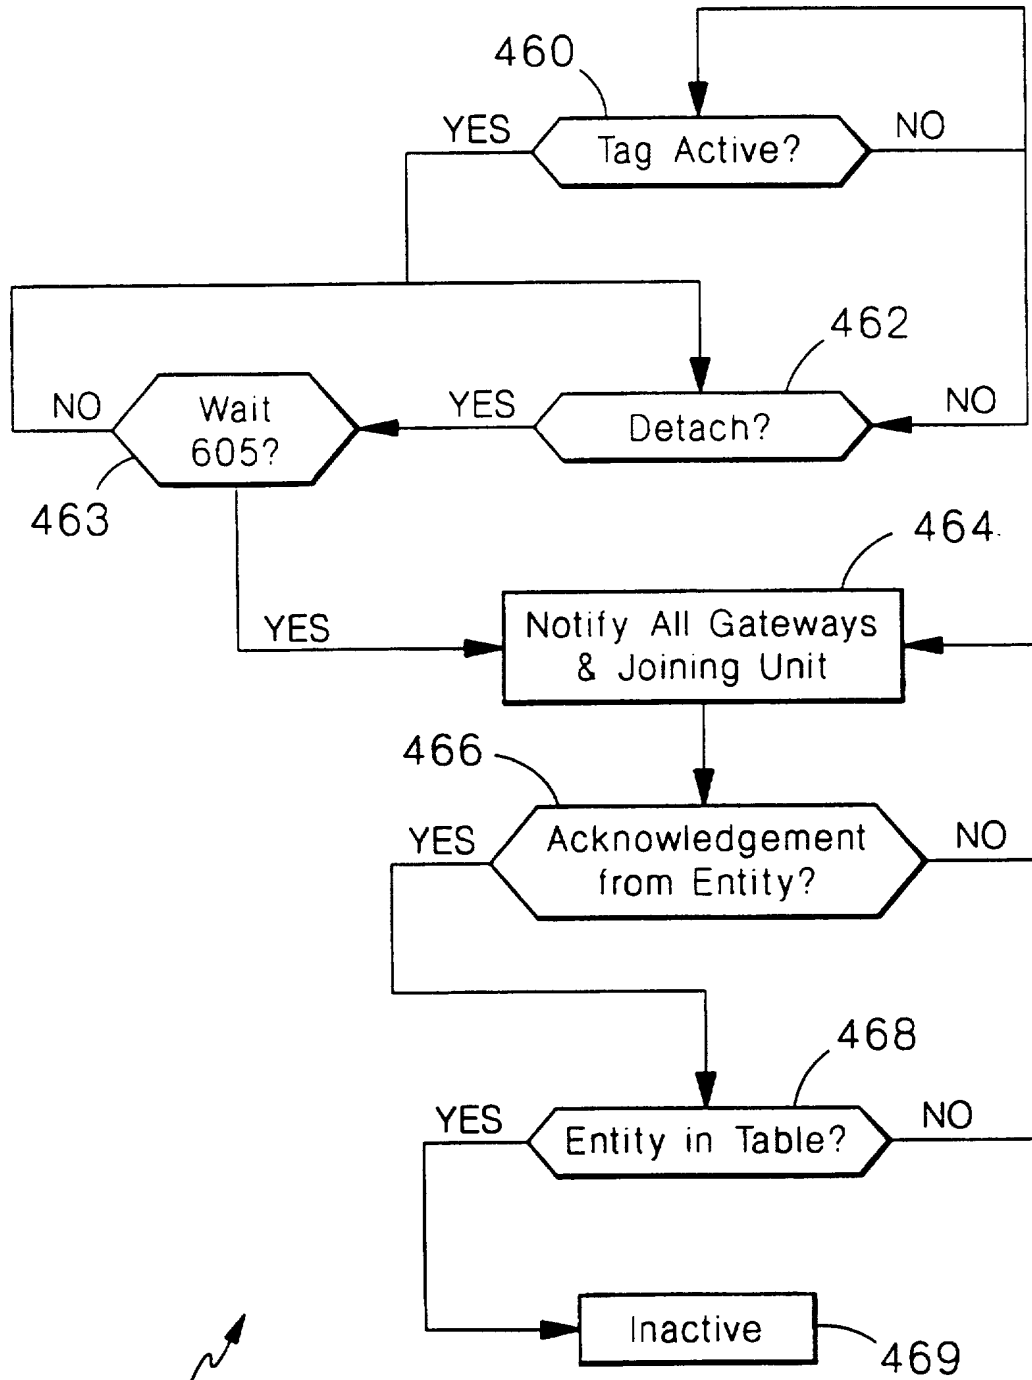



FIG. 27



404 ↗

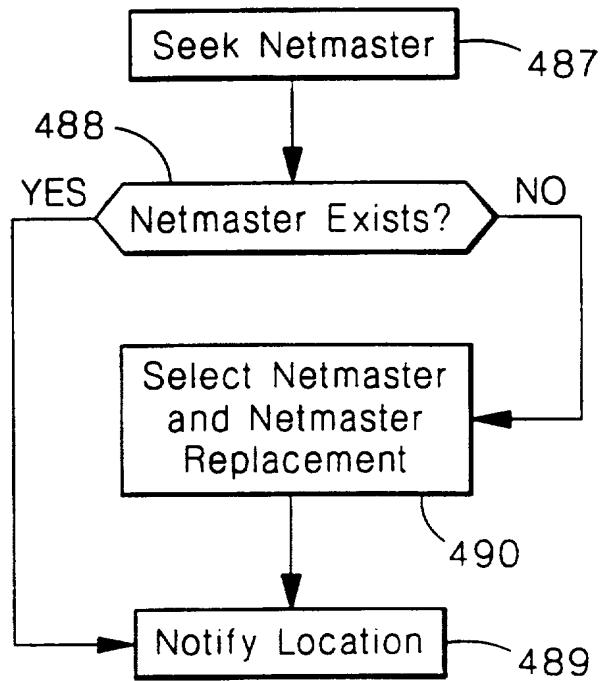
FIG. 28



470 ↗

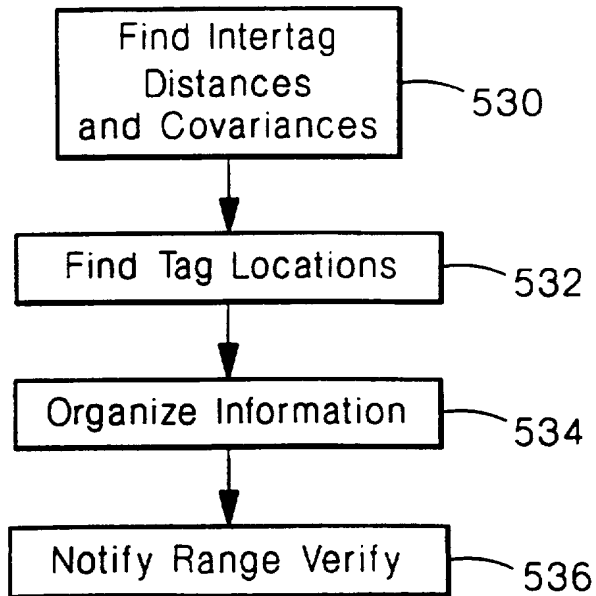
FIG. 29

30/59



472 ↗

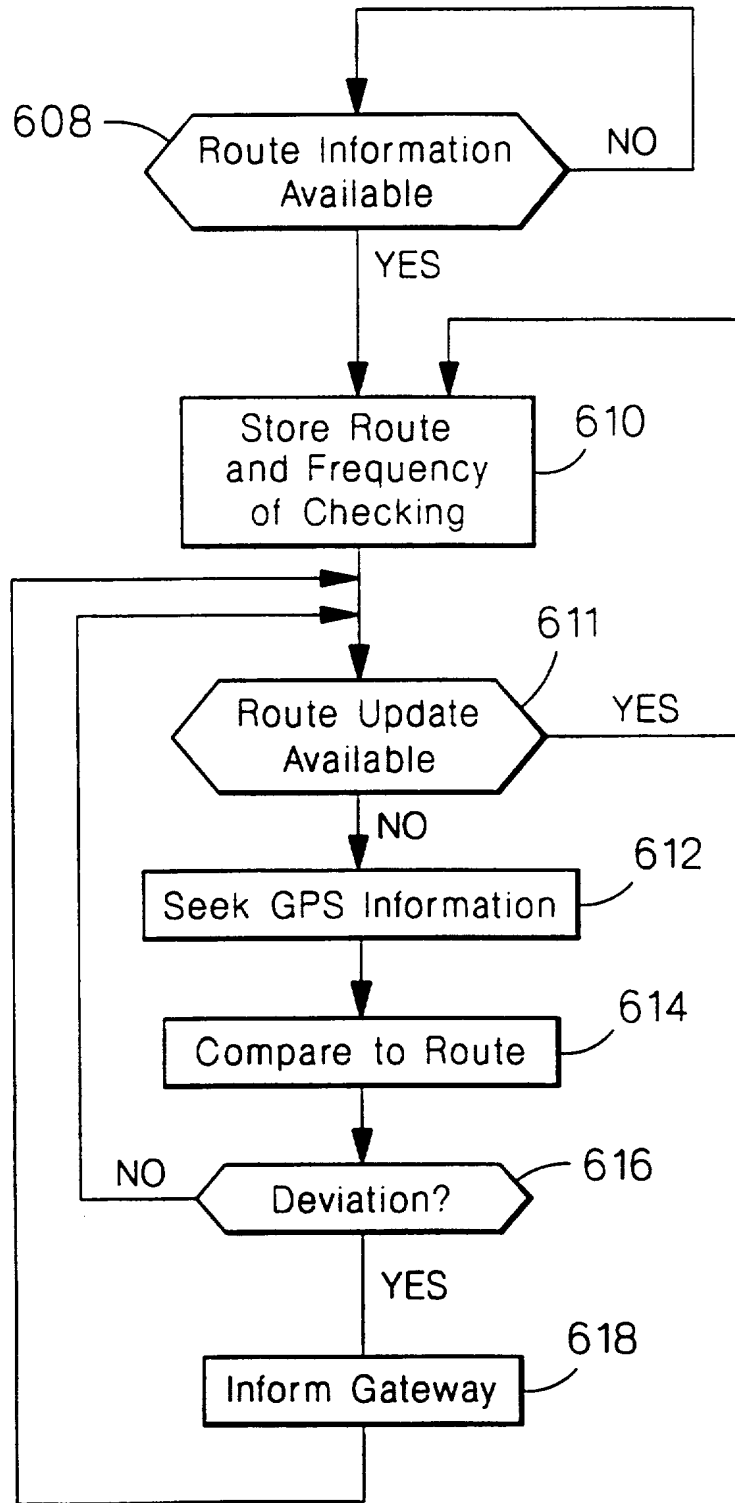
FIG. 30



500 ↗

FIG. 31

31/59



590 ↗

FIG. 32

32/59




FIG. 33

592 ↗

Transponder Tag



Range



Originator Tag

FIG. 34




FIG. 36

35/59

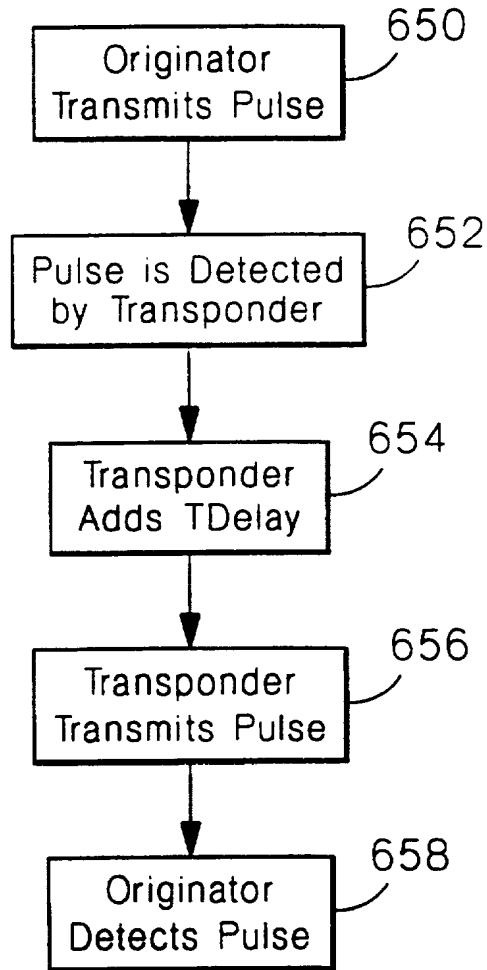


FIG. 37

36/59




FIG. 38




FIG. 39

37/59

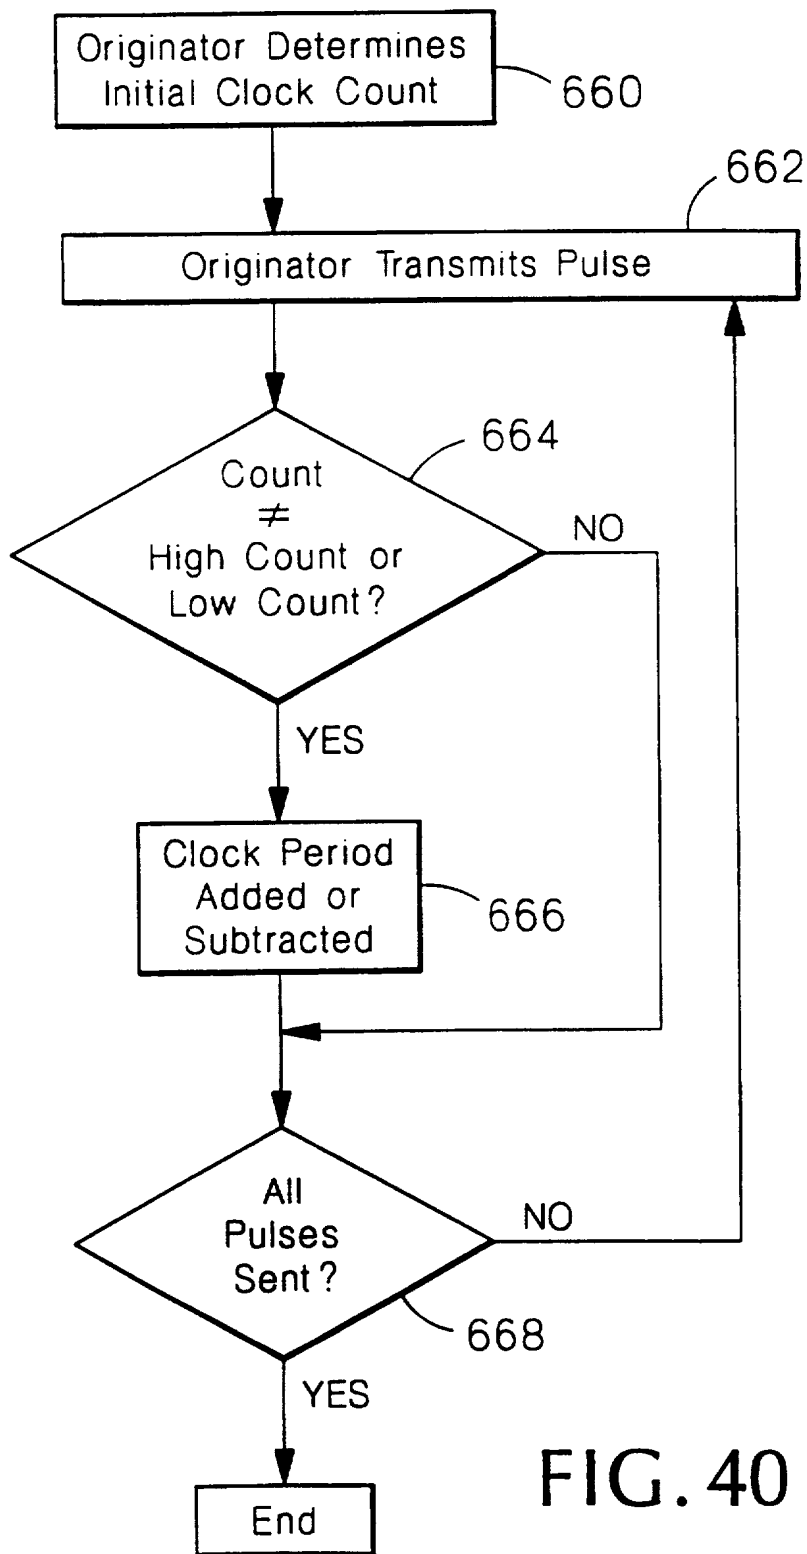


FIG. 40




FIG. 41

39/59

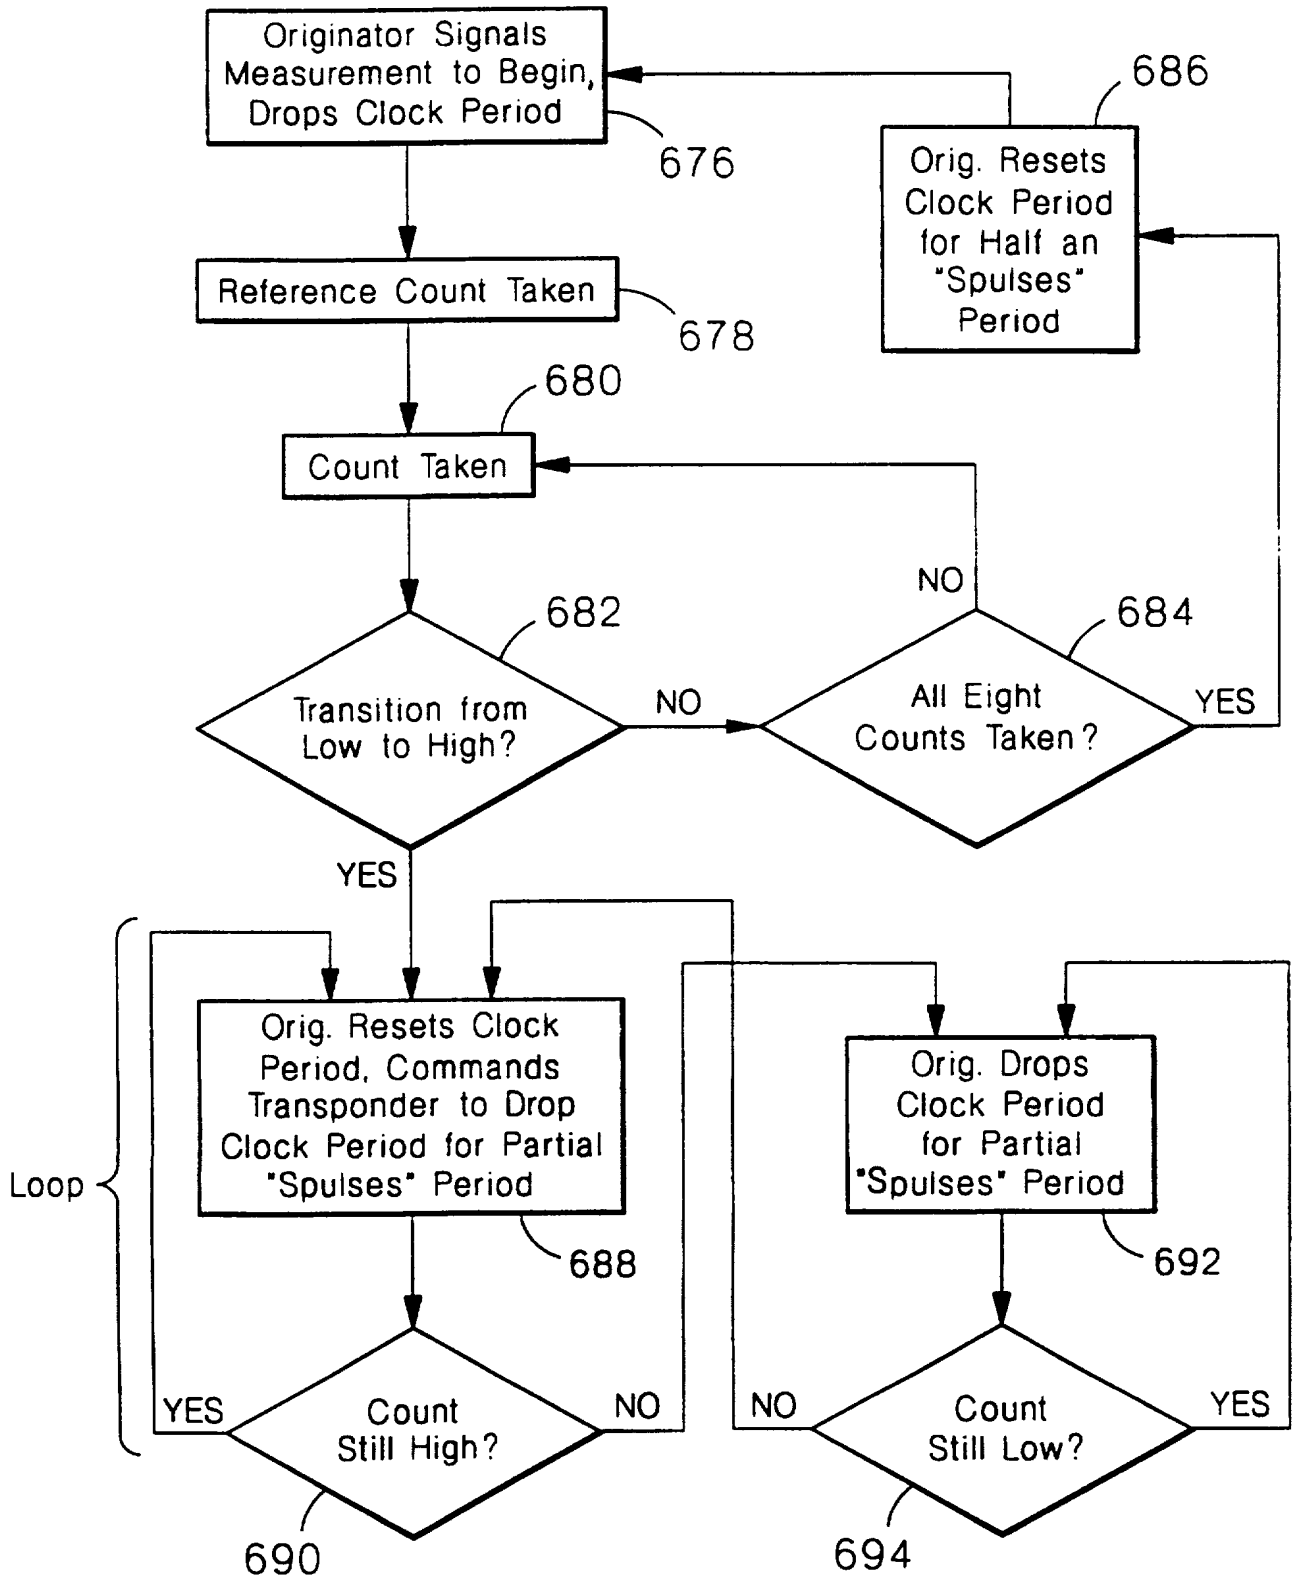


FIG. 42

41/59

800

Pair	Range	Range σ	Number of Errors
Pair 1	f_1	σ_1	e_1
Pair 2	f_2	σ_2	e_2
etc.			

FIG. 44

820

Triplet Error	Pair 1	Pair 2	Pair 3
Triplet Error 1			
Triplet Error 2			
etc.			

FIG. 47

820'

4-Tuple Error	Pair 1	Pair 2	Pair 3	Pair 4	Pair 5	Pair 6
4-Tuple Error 1						
4-Tuple Error 2						
etc.						

FIG. 49

820''

5-Tuple Error	Pair 1	Pair 2	Pair 3	Pair 4	...	Pair 10
5-Tuple Error 1						
5-Tuple Error 2						
etc.						

FIG. 50

42/59

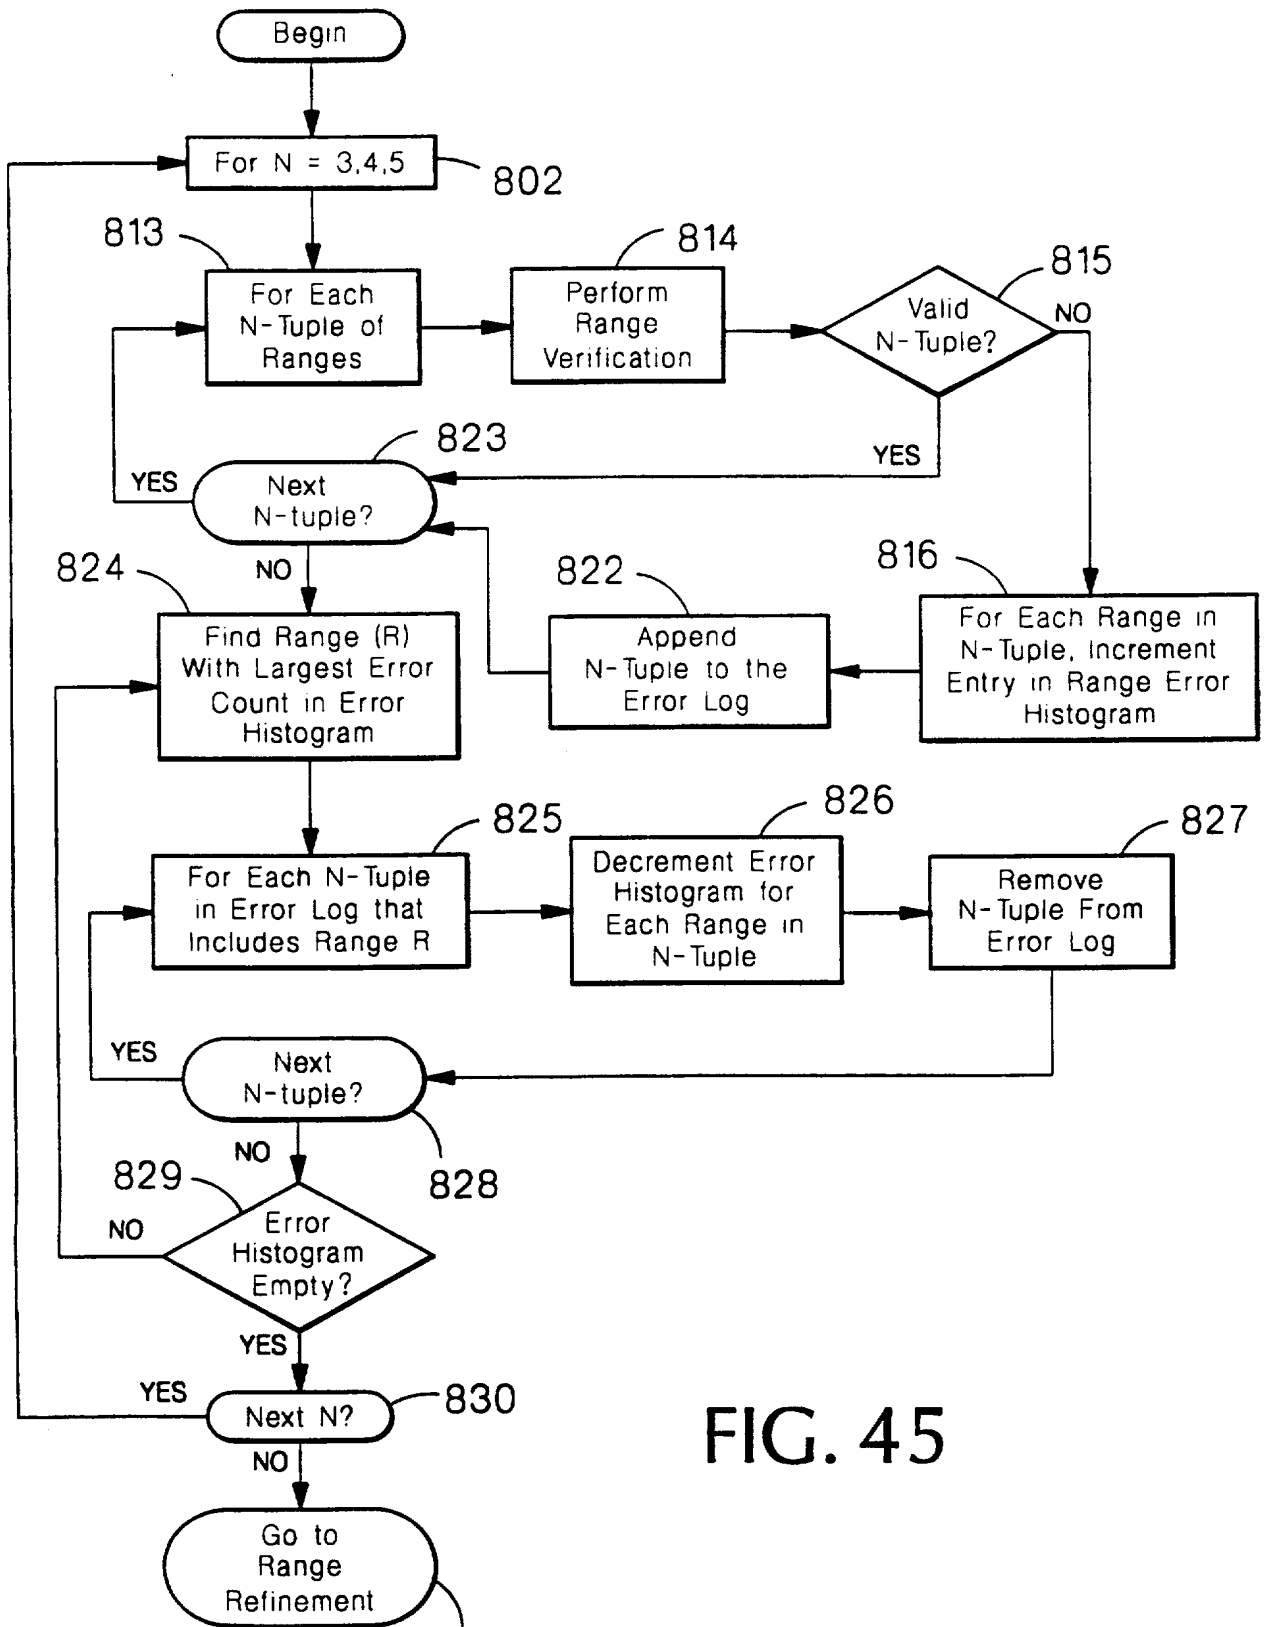


FIG. 45

830a

43/59




FIG. 46




FIG. 48

44/59




FIG. 51

45/59

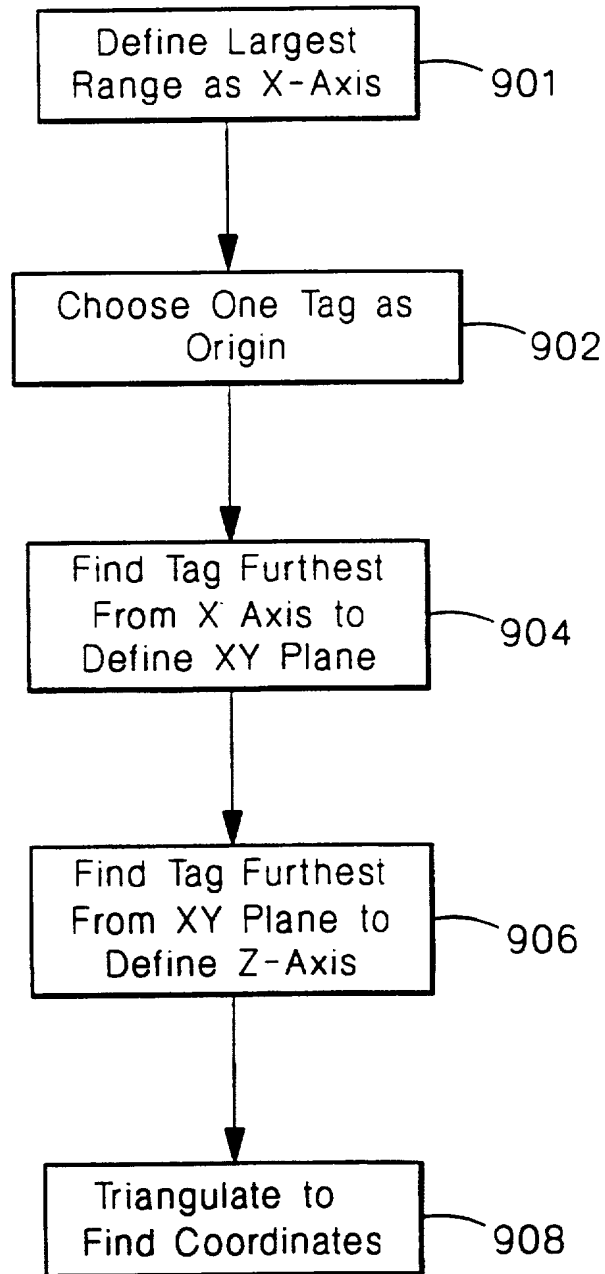


FIG. 52

46/59




FIG. 53




FIG. 54

47/59

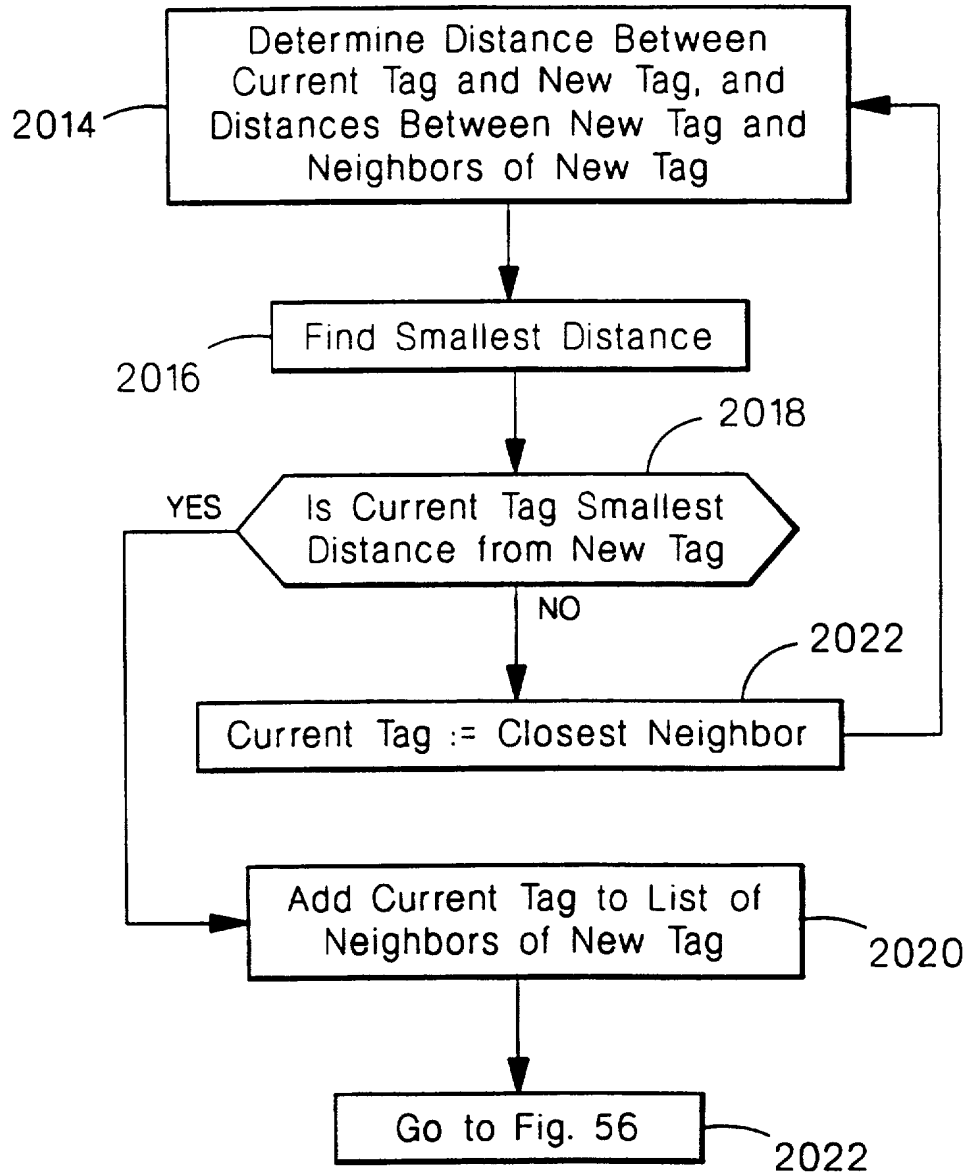


FIG. 55

48/59

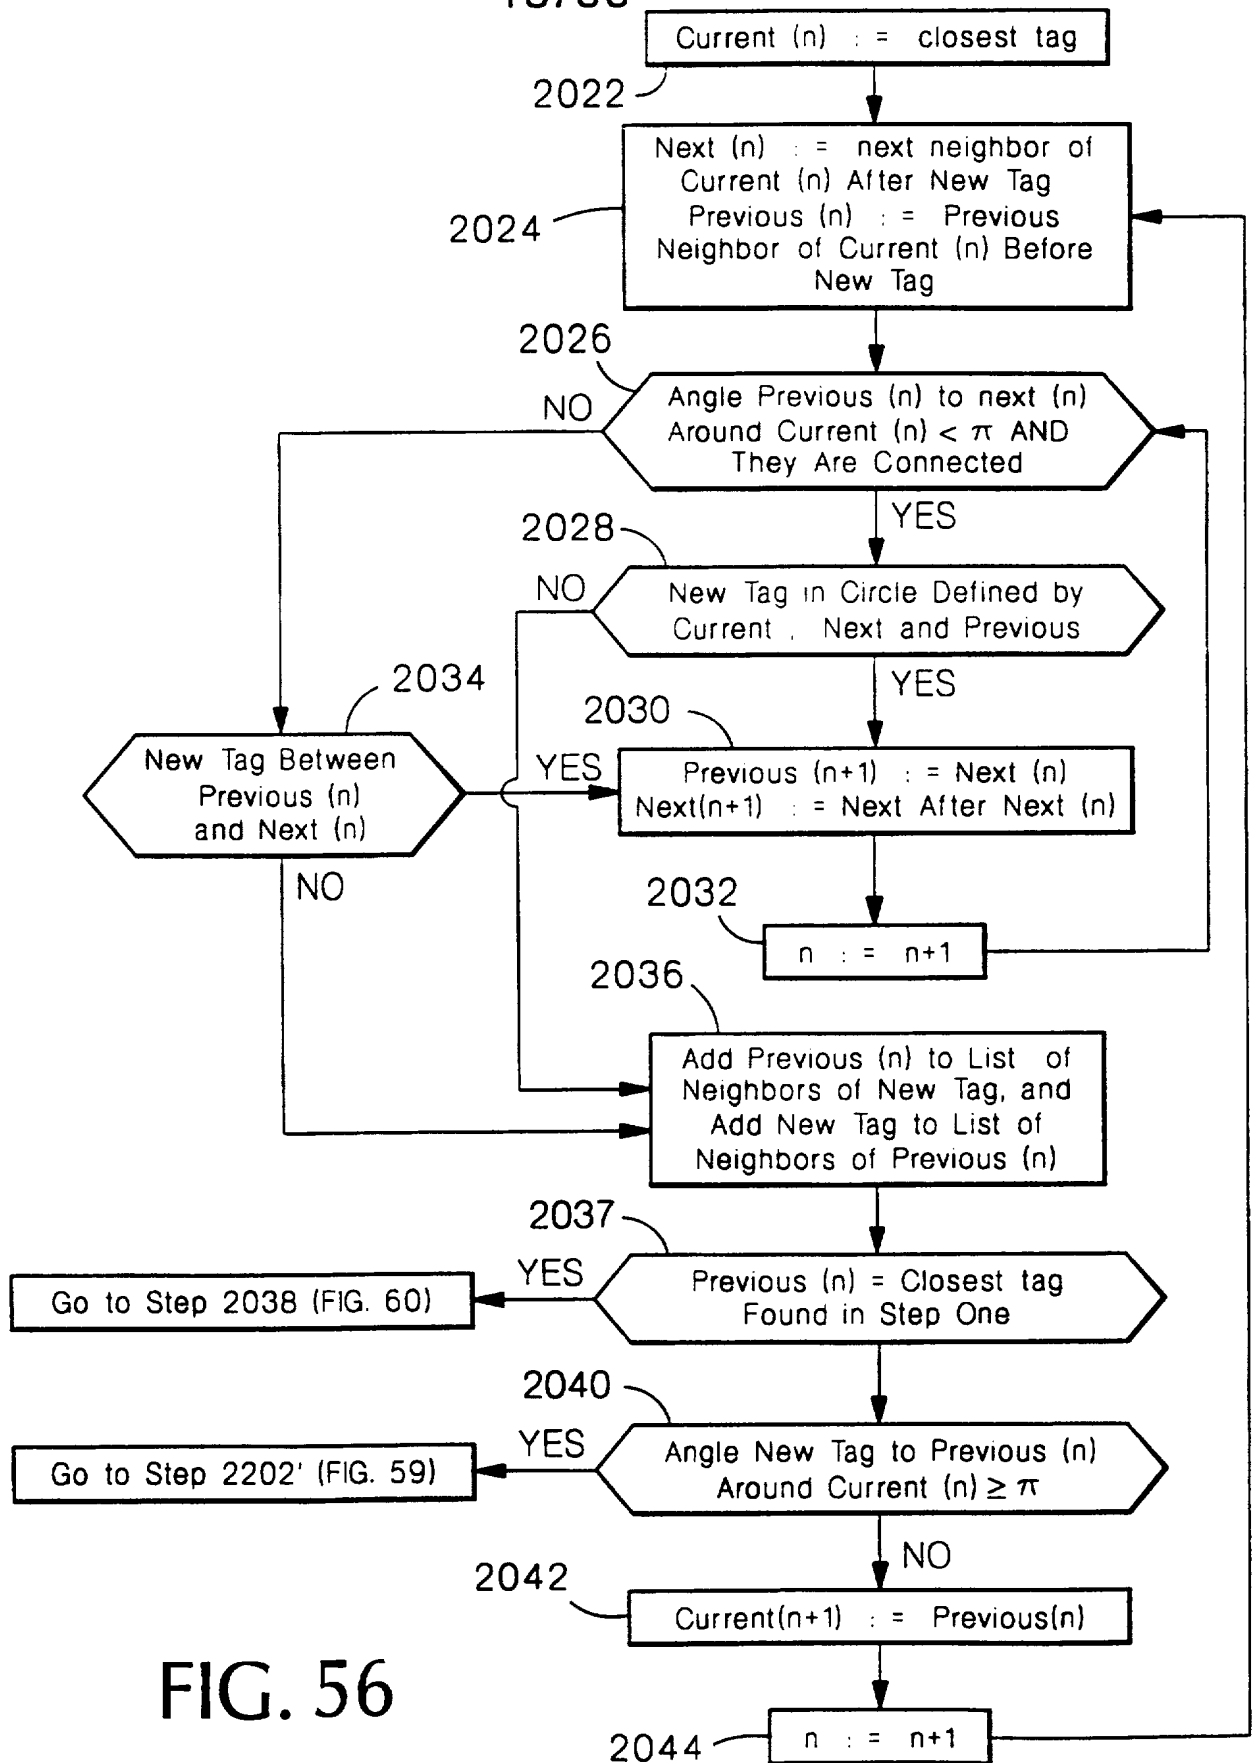


FIG. 56




FIG. 57

50/59




FIG. 58




FIG. 61

SUBSTITUTE SHEET (RULE 26)

51/59

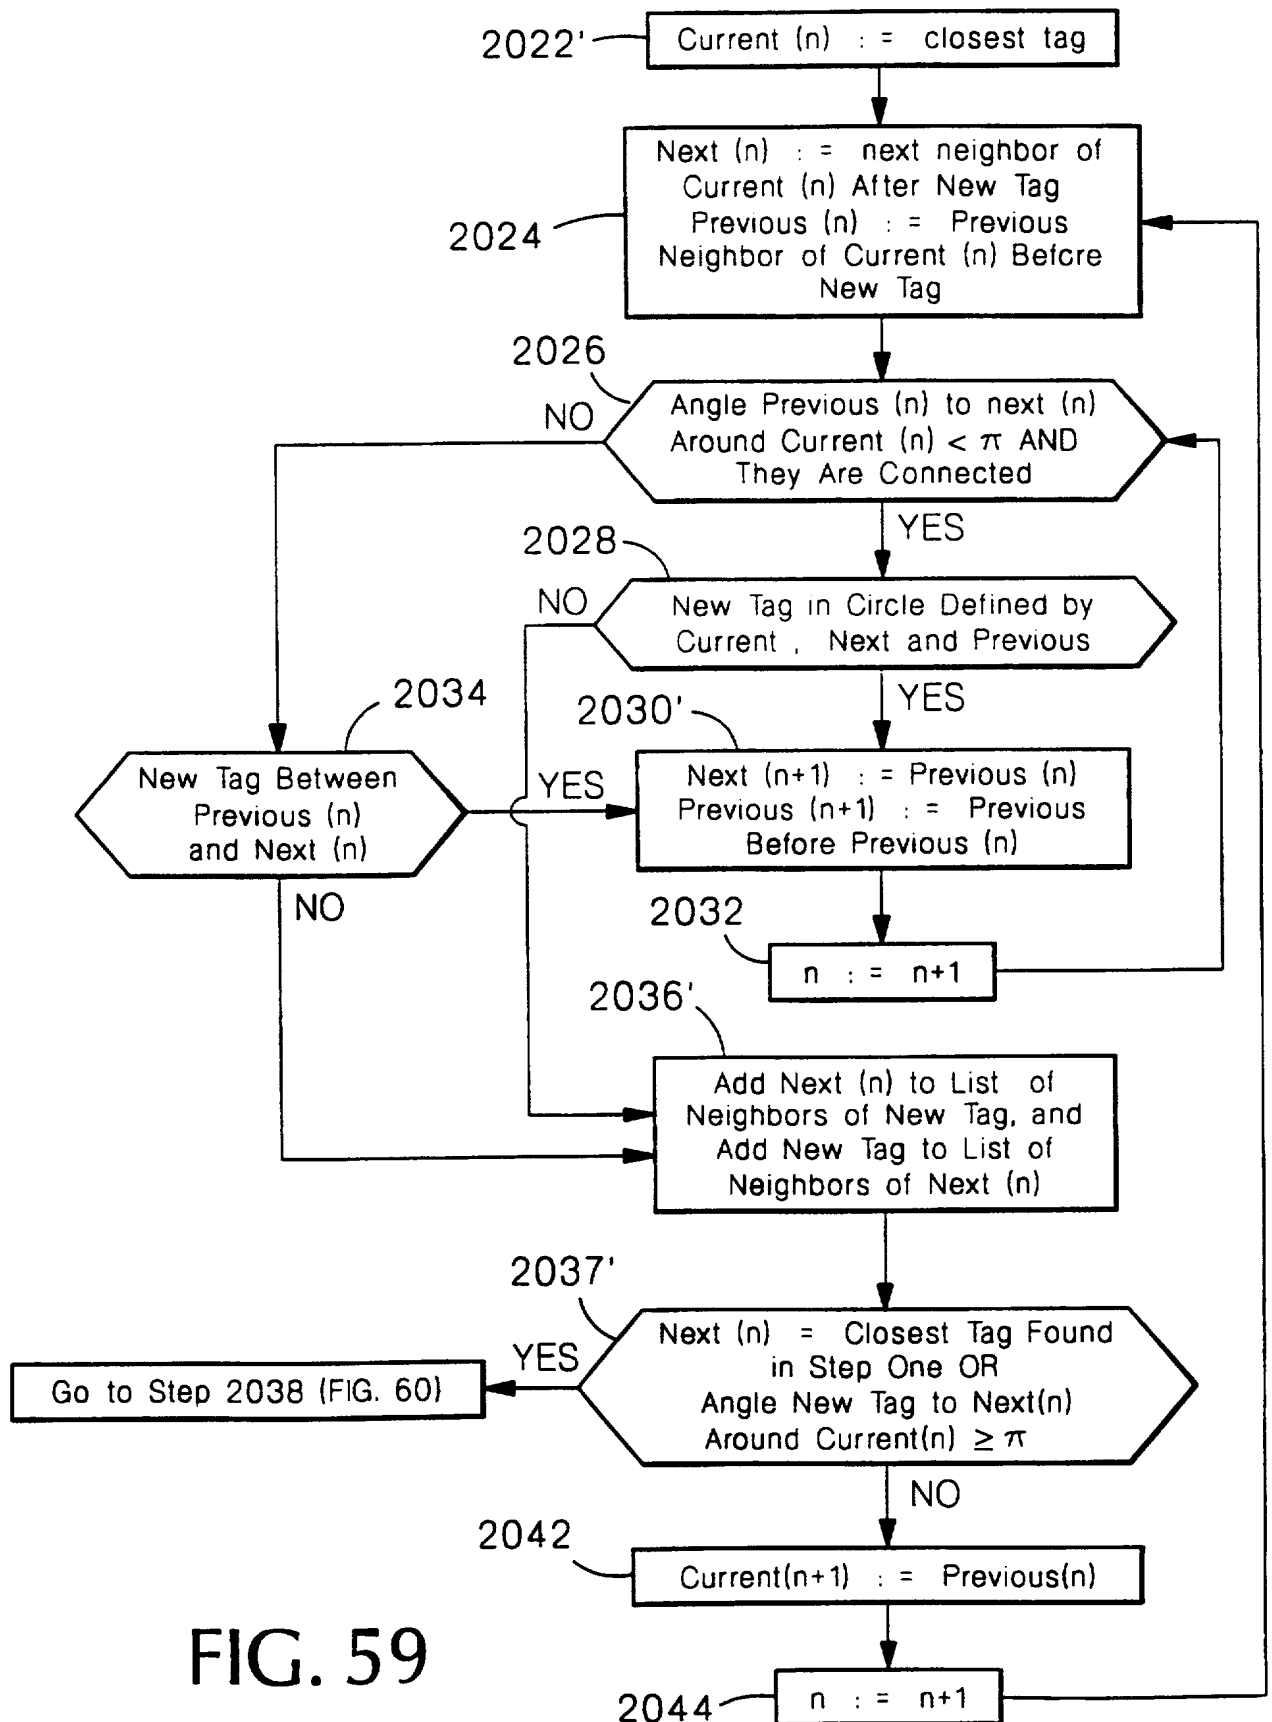


FIG. 59

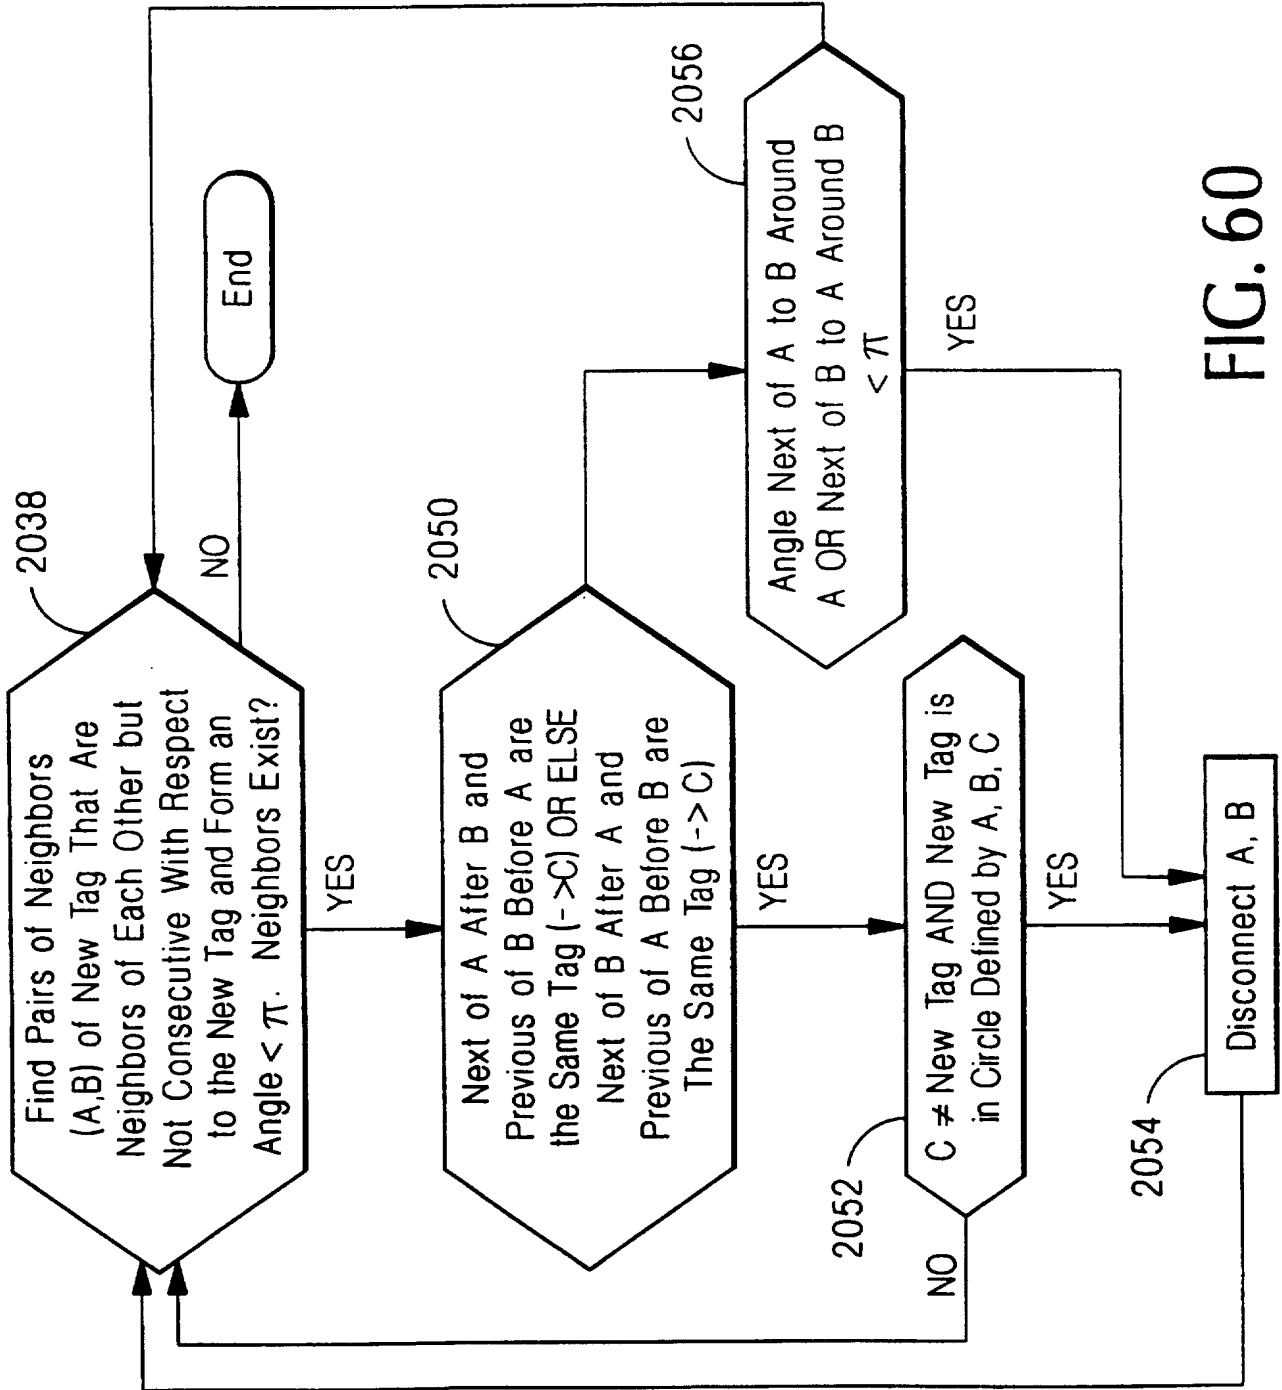


FIG. 60

54/59




FIG. 63

104

55/59




FIG. 64

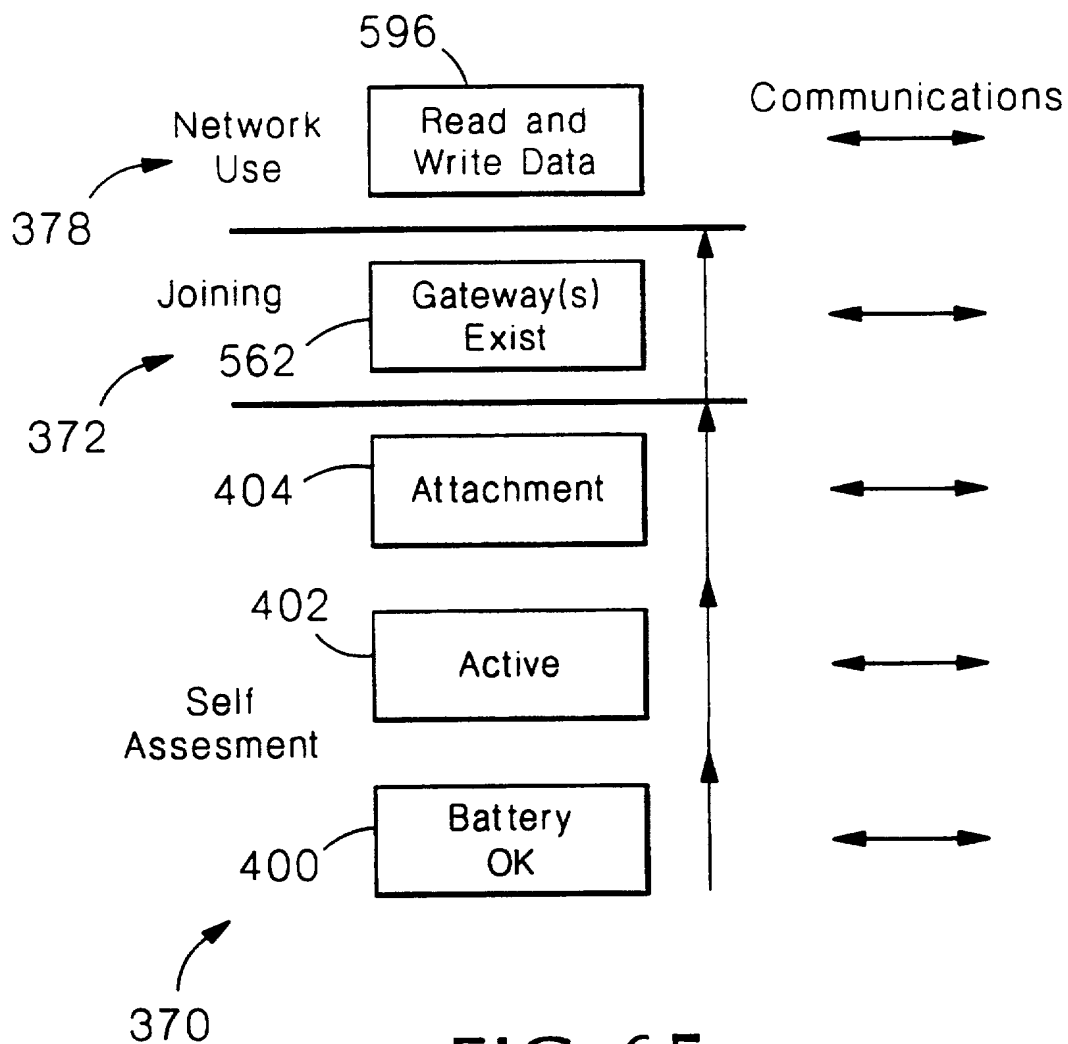


FIG. 65

57/59




FIG. 66

58/59




FIG. 67





FIG. 68

SUBSTITUTE SHEET (RULE 26)



59/59

FIG. 69